Millimeter-Wave Full Duplex Radios

Vaibhav Singh
Carnegie Mellon University
vaibhav3@andrew.cmu.edu

Susnata Mondal
Carnegie Mellon University
susnatam@andrew.cmu.edu

Akshay Gadre
Carnegie Mellon University
agadre@andrew.cmu.edu

Milind Srivastava
Indian Institute of Technology Madras
milind@cse.iitm.ac.in

Jeyanandh Paramesh
Carnegie Mellon University
paramesh@andrew.cmu.edu

Swarun Kumar
Carnegie Mellon University
swarun@cmu.edu

ABSTRACT

mm-Wave has emerged as an attractive high-speed wireless communication paradigm owing to the high available bandwidth at mm-wave frequencies. Full-Duplex has the potential to double the available capacity in the mm-wave bands by enabling simultaneous radio transmission and reception. While full-duplex has been extensively studied in sub-6 GHz bands, this paper exposes the unique challenges in porting this capability to mm-wave frequencies.

We present mmFD, the first comprehensive system design of a mm-wave full-duplex platform. mmFD achieves large self-interference cancellation through novel designs at the antenna, analog and digital frontends. We exploit the small wavelength of mm-wave to achieve strong signal isolation between transmit and receive antennas. We further build a custom IC that achieves high-bandwidth analog cancellation at mm-wave frequencies. Finally, we present digital self-interference cancellation algorithms that address the unique hardware impairments observed at mm-wave frequencies. A detailed evaluation of mmFD demonstrates 84 dB of cancellation and 1.7× throughput gain over equivalent half-duplex systems in rich indoor settings.

CCS CONCEPTS

• **Networks** → **Network components**; • **Computer systems organization** → **Embedded and cyber-physical systems**; • **Hardware** → **Integrated circuits**.

KEYWORDS

millimeter wave, self-interference cancellation, RFIC, self-reflector

ACM Reference Format:

Vaibhav Singh, Susnata Mondal, Akshay Gadre, Milind Srivastava, Jeyanandh Paramesh, and Swarun Kumar. 2020. Millimeter-Wave Full Duplex Radios. In *The 26th Annual International Conference on Mobile Computing and Networking (MobiCom '20)*, September 21–25, 2020, London, United Kingdom. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3372224.3380879>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

MobiCom '20, September 21–25, 2020, London, United Kingdom

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-7085-1/20/09...\$15.00

<https://doi.org/10.1145/3372224.3380879>

1 INTRODUCTION

Recent years have seen major advances in millimeter wave technology to push the capacity of wireless links to cater to the ever increasing data needs in both the cellular and wireless local area network scenarios. While mm-wave suffers from a high path loss, it is a good candidate for delivering high speed data traffic for short range applications like WiFi and small-cell cellular due to its large available bandwidth. However, with mobile data traffic slated to increase by 53% each year [10], it is imperative to ensure that mm-Wave systems offer the maximum spectral efficiency possible.

Full duplex is an attractive solution in this direction, with the potential to increase the network capacity nearly two-fold. It achieves this by transmitting and receiving at the same time on the same frequency channel unlike FDD and TDD systems where transmission and reception occur on different frequency/time blocks respectively.

Figure 1: mmFD: Bidirectional full-duplex mm-Wave

Full duplex communication has been studied extensively in the last decade, primarily targeting sub-6 GHz frequencies. The main challenge associated with any full duplex system is the self interference (SI) caused by the transmitter at its own local receiver. Past systems employ a combination of interference cancellation techniques[35] in the analog and digital domain to overcome these challenges. Recent years have seen major advances in full-duplex design for the 2.4 GHz Wi-Fi frequencies [9, 15] as well as associated MAC protocols [31]. There have also been theoretical models for full duplex in the mm-wave context for different antenna placements [71]. Perhaps the closest attempt to demonstrating a full-duplex link appears in [18], which demonstrates a full duplex link at 60 GHz using polarization based antennas and a CMOS IC, but they use the same clock for both the nodes of the link, thus eliminating

dealing with synchronization related issues. Another close attempt is [59], which uses two adjacent, well-separated panels of two large base stations as the transmitters and receivers to demonstrate concurrent bi-directional communication at 28 GHz. However, there is no concrete system today that deals with the problem of full duplex for mm-Wave consumer devices. Thus, there remains a need for enabling full-duplex communication between small form-factor mm-Wave end-user devices or access points.

This paper presents mmFD, the first comprehensive study of the unique challenges of a bi-directional mm-Wave link in the full duplex context, compared to the sub-6 GHz. We present the first SDR-driven study for developing solutions to overcome these challenges by using a combination of isolation/cancellation techniques at the antenna front-end, analog and digital domain. We further evaluated a mm-Wave full-duplex capable integrated chip designed in-house[52]. mmFD achieves an average $1.7\times$ increase in throughput compared to the half duplex system at 27.35 GHz. We achieve this through cancellation at three stages in the full-duplex signal path: the antenna, the analog frontend and digital frontend. Below we describe the unique challenges in the mm-wave context in achieving this cancellation across these stages.

Antenna Design: The first challenge associated with mmFD is the high frequency self-interference signal at the antenna front end. Traditional sub-6 GHz full-duplex systems, typically, employ either large antenna separation, cross polarized antennas or circulators to achieve a large degree of isolation between the receiver and transmitter lines. In contrast, mm-wave allows for a more elegant antenna isolation solution, given that its wavelength is relatively small. We design a small customized self-reflector co-located with transmit and receive patch antennas (controlled independently) capable of generating high cancellation at the receiver antenna without significantly affecting the desired signal. We describe why our approach is well-suited for mm-wave clients with small form factors (described in Sec. 4). Our system achieves 30 dB + 10 dB cancellation owing to antenna isolation and self-reflector cancellation respectively at the antenna front.

Cancellation in Analog: Despite some degree of antenna isolation, there remains considerable self-interference in the received signal due to residual interference at the antenna as well as the spill-over interference due to the proximity of the transmit and receive lines within the chip. To cancel the self-interference in the analog domain, traditional full-duplex systems apply a cancellation signal that matches and negates self-interference. A key challenge in the mm-wave context is to generate the cancellation signal quickly at the high bandwidth of mm-wave. We achieve this through the design of a custom mmFD chip which passes a copy of the transmitted signal through a fine gain and phase control block to adaptively generate a cancellation signal that matches self-interference. Our innovation lies in optimizing the weights of our chip without exhaustively evaluating over 2^{20} possibilities. We present our solution as well as how these weights can be adapted over time in Sec. 5. We demonstrate average analog cancellation of 18 dB over a bandwidth of 100 MHz at 27.35 GHz.

Cancellation in Digital: Our above optimization reduces the power of self-interference signal sampled at the digital front-end, enough to avoid saturation at the Analog-to-Digital Converter (ADC). Yet

there is still residual interference signal at the digital front-end. While the transmitter and receiver at one end operate using the same clock removing any frequency offsets, even small subsample timing offsets create large variations in phase (more than π) across the wide bandwidth. Thus, using traditional one-tap or multi-tap cancellation leads to amplification of interference on outer bands of the wideband signal. Further, time-variant distortions caused due to high-frequency components need to be addressed. We present a novel approach that performs wide-band phase-sensitive digital cancellation even at low self-interference-to-noise ratio and eliminates the time-invariant distortions (described in Sec. 6). Our digital cancellation shows improvement in SINR of 26 dB on average.

Evaluation: We implement our mmFD node using the custom chip as the RF front end with patch antennas. We design a custom TRX board to up and downconvert the signal to 27.35 GHz. USRP X310 is used to generate the baseband signal. We use MATLAB for digital processing of the baseband signal for self-interference cancellation. Our results reveal:

- A total self-interference cancellation of 84 dB for a bandwidth of 100 MHz
- A throughput gain of $1.7\times$ over traditional mm-Wave half-duplex systems.
- mmFD's cancellation leaves less than 1 dB residual interference in multipath-rich indoor environments across 8m, both LOS and NLOS, at 0 dBm transmit power.

Contributions: This paper's main contributions are:

- The first SDR-driven comprehensive study of a bi-directional full-duplex communication system in the mm-wave context.
- A system design that addresses the mm-Wave specific challenges in eliminating full-duplex self-interference via antenna design, analog and digital frontends.
- Evaluation of a custom IC to perform full-duplex analog cancellation at mm-wave frequencies in multipath-rich indoor settings.

2 RELATED WORK

The related works can be broadly classified as follows:

Millimeter Wave Protocols and Applications: With the emergence of 5G and mm-Wave technology, there has been a lot of work on designing systems with large bandwidth, smaller elements size and narrow beams offered at these frequencies to achieve larger throughput and accurate sensing. Much of the research in this domain has focused on developing novel antenna array designs [25, 30, 55, 58, 72], and fast and efficient beam training and tracking [4, 28, 37, 46, 48, 53, 60, 67, 68, 70]. Other approaches [5, 20, 29, 30, 33, 47, 49–51, 54] have explored combining massive arrays of antennas in MIMO configuration to improve the range and throughput of mm-Wave systems. In the circuits community, efficient mm-Wave circuit and antenna design to combat high frequency non-linear distortions and phase noise have been explored in [22, 38, 40, 64]. 802.11ad/ay [24, 57] are the IEEE standards for 60 GHz WiFi with well-defined network stacks and associated MAC protocols. mm-Wave radar[36] and imaging sensors[74] have been

designed to aid in localization and vehicular automation applications. Our work builds upon the knowledge from these systems to come up with a full duplex system at these frequencies.

Full Duplex: Full Duplex is a well studied problem with different systems employing a combination of antenna isolation and analog/digital cancellation techniques. Some solutions [7–9] present full duplex WiFi system operating in the 802.11ac's 2.4 GHz band and later builds up on that for a MIMO full duplex system achieving 90 dB SIC. Other papers [17, 23, 31, 41, 45] present MAC layer designs for WiFi full duplex systems which evaluate the improvement achieved in terms of network performance. [7, 16, 26, 61, 62] used full duplex for wireless relays but these essentially involve interference nulling with a combination of analog and digital domain SIC for the residual interference as well as beamforming based nulling using antenna arrays. [9, 15, 19, 31] use complex weight scaled time delayed versions of transmit signal for interference cancellation whereas [11, 13, 73] frequency domain equalization for wideband SIC. However, the works discussed above deals with sub-6 GHz frequency and hence are not applicable for mm-wave.

mm-Wave Full Duplex: In the context of full duplex for mm-Wave systems, there have been extensive theoretical studies on using beamforming cancellation techniques [1, 27, 39, 44, 56, 63, 71] in addition to the RF and digital cancellation for backhaul and relay networks, however there have been no system implementation based on these. mm-Wave FMCW automotive radars employ analog cancellation techniques [21, 34, 42] to suppress the self-interference. However, mm-Wave radars pose less challenging demands than communication devices since they do not have to deal with SIC of wideband signals with complex modulations, and furthermore employ a single clock due to their monostatic nature. Recent years have also seen interesting combination of mm-Wave and fiber optic communication technologies [2, 6, 12, 32, 43] to create a full duplex system by modulating a mm-wave signal over a optical frequency to be sent over an optical fiber which leverage the orthogonal polarizations of the light signal to achieve full duplex communications.

Perhaps the closest work to ours are solutions [14, 18] which test a 60 GHz custom chip in a uni-directional full-duplex scenario. Yet, both solutions use polarization to reduce Self-Interference (SI) which need not operate in complex multipath scenarios. Moreover, [14] is particularly expensive in terms of area and power consumption, and is not easily amenable to scaling to multiple streams. Thus, all prior work fails to delineate the challenges associated with the system and rather focuses on describing the cancellation and the associated SI reduction. In contrast, our work describes the hardware challenges in deploying a mm-Wave full duplex system as well as does the first comprehensive evaluation of a bidirectional full duplex link with different clocks and real-world PHY-layers in multipath rich indoor settings.

Our work builds upon prior design of a hybrid beamforming custom chip [52] that can be configured to operate in full duplex mode. The prior work is a short paper which highlights its modular nature enabling different modes of configuration at scale: MIMO TX, MIMO RX or full-duplex. It characterizes the performance of its circuit blocks in terms of power consumption, TX-RX configurability, etc. as well as cancellation on a wafer probe station. mmFD presents the first over-the-air evaluation of the chip, optimizing its

parameters for cancellation in real-time (See Sec. 5). Further, mmFD deals with the other system level challenges in demonstrating a bidirectional full duplex link: interfacing with commodity SDRs via custom-designed TRX board, developing antenna cancellation solution for patch antennas in mm-Wave setting and cancelling residual interference in software.

3 FULL DUPLEX: MM-WAVE VS. 2.4 GHZ

In this section, we explore how the full-duplex problem changes from sub-6 GHz to 28 GHz. Much like traditional full-duplex, our approach cancels self-interference from the transmit to the receive chain on the full-duplex node at three fronts: the antenna, analog and digital. Below, we summarize the challenges and opportunities enabled by mm-wave in maximizing cancellation in each of these components.

Antenna Isolation: Traditionally, sub-6 GHz full-duplex systems use circulators to isolate the transmit and receive streams using single antenna or rely on careful spacing of antennas to isolate interference between transmit and receive antennas. However, circulator based systems are ill-suited at mm-wave for two reasons: size and inability to scale to multi-antenna systems. Indeed, mm-wave full-duplex systems can take a similar approach in achieving some degree of antenna isolation. However, we note that the small wavelength and high bandwidth of mm-wave affords a new opportunity to cancel self-interference: *self-reflectors*. Specifically, we seek to engineer multipath reflections that cancel out much of the dominant components of self-interference. We do this by carefully placing thin metallic strip in the vicinity of the transmitter and receiver that introduces additional signal paths that precisely cancel self-interference to a high-degree. Such an approach would be unwieldy and fickle in the sub-6 GHz domain, as the reflector would have to be large and far away from the full-duplex node to match self-interference – increasing the node's form factor. Sec. 4 details our approach and discusses means to optimally position this reflector, without impacting the desired signal.

Analog Cancellation: While our self-reflector coupled with antenna isolation provides significant cancellation, the residual interference is still about 40-45 dB above the noise floor. Further, significant interference leaks into the RX chain at the antenna front as well as within the chip. We use a custom chip at the analog front to cancel a significant portion of this interference by modelling the residual self-interference channel, to allow the SDR ADC to capture the desired signal. Unlike previous mm-Wave full-duplex solutions which favor single antenna systems, our self-interference cancellation chip can seamlessly scale to multiple antennas. The self-interference cancellation signal is a weighted copy of the transmitted signal. A major challenge for our system to work at high bandwidths is to quickly learn and update this weight to match the self-interference component by estimating the channel of the leakage path. We describe our chip-design and fast-update algorithm in Sec. 5.

Digital Cancellation: Finally, we use digital approach to cancel out any residual interference. Prior approaches, particularly at sub-6 GHz have dealt with such cancellation by modelling filters, hardware offsets and non-linearities in both time-domain and

Figure 2: mmFD's antenna design: (a) Self-reflector is a thin strip of metal which upon optimum placement cancels the interference signal at the receive antenna. (b) Coupling power decreases drastically as distance between antennas, d , is increased. (c) Variation of SINR across locations of self-reflector, l , shows a maxima within λ .

frequency-domain. At the high center frequency and wide bandwidth of mm-Wave, these effects are significantly more exacerbated. Further, high frequency components create larger time-invariant distortions due to saturation of hardware components. We develop an adaptive approach that compensates interference for each frequency and time based on the observed phase thus making the system robust to frequency dependent distortions of residual self-interference channel after analog cancellation. We present the above approach and how to eliminate time-invariant distortions in Sec. 6.

4 MMFD ANTENNA DESIGN

In this section, we describe in detail our approach to mitigate self-interference between the two antennas of the full-duplex node. We want to minimize leakage of signal across antennas while maximizing the throughput of the participating links in the full-duplex system.

Given the highly directional nature of mm-Wave beams, distance between the transmit and receive antennas would be enough to mitigate most of the interference. Our experimental evaluation (shown in Fig. 2(b)), shows that the interference indeed reduces as we move our patch antennas further distance (d) away. Yet, doing so would be at the expense of increasing the size of full-duplex mm-Wave antenna chips.

Previous solutions in the WiFi context have explored ways to null its own signals at the Rx antenna using circulators or used efficient placement of antennas to minimize interference. However, circulator based systems are ill-suited for mm-Wave application for two reasons: First, they are bulky to implement on chip in a small form factor which results in them occupying a large die area (there have been some implementation of on-chip circulators, but their performance is not satisfactory). Second, they are limited in their ability to deal with interference only from their own TX element but not from adjacent TX elements as is the case in multi-element phased arrays.

Further, mm-Wave solutions have leveraged orthogonality of polarization to minimize coupling and maximize the signal-to-interference-noise ratio (SINR) at the receiver. However, these solutions still suffer from coupling interference caused on the chip. The heavy attenuation over the air means that the SINR is also significantly worse over the small distances. Luckily for us, mm-wave frequencies offer an interesting opportunity - a small wavelength.

4.1 Self-Reflector for Cancellation

We present the idea of a *self-reflector*, a thin strip of metal deployed on the axis perpendicular to that of TX and RX antennas to cancel out the leaking interference signal. Note, our aim is only to reduce the effect of the most dominant tap by engineering the multipath of the interference channel. Furthermore, our evaluation and recent studies [66, 69] have shown that due to the high directionality of phased arrays and heavy attenuation of mm-Wave signal over the air, power of the most dominant tap is significantly larger than the other taps. Our approach stems from the above analysis which shows that the dominant tap of the interference signal received at the receiver antenna is, $r[t] = Ae^{i\theta}x[t]$, an attenuated and phase shifted version of the transmitted signal. Our approach suggests sending another signal (reflected from our *self-reflector*) from transmit antenna to the receive antenna over the air whose phase is opposite.

$$y[t] = \frac{1}{k}e^{i2\pi\frac{k}{\lambda}}x[t]$$

This signal will cancel the interference at the receiver antenna, thereby improving the SINR. Note that if this path is too long, the reflector will have negligible effect on the interferer signal.

Our self-reflector is placed on a fine-grained movement platform which can be moved to change the phase and attenuation of the reflected path at the receiver antenna. Fig. 2(c) shows how across distance the attenuation increases and decreases as the phase of reflected signal changes. We can see that the effect is periodic as well as it reduces as the reflector is placed further away. Furthermore, our results also demonstrate that such a reflector will not significantly affect the legitimate signal coming from the other client. The above evaluation shows that correct placement of the reflector can improve SINR by upto 10 dB at a small form factor.

For the self-reflector, the distance over the air of the reflected signal can be characterized as $k=2\sqrt{\frac{d^2}{4} + l^2}$ where l is the distance of the self-reflector from the antenna plane and d is the distance between antennas. Thus, to maximize cancellation our objective function becomes

$$\min_l \Phi = \|Ae^{i\theta} + \frac{1}{k}e^{i2\pi\frac{k}{\lambda}}\|_2^2$$

Note that the power of sum of two complex numbers is minimized when they are perfectly out of phase, that is, the phase between the

Figure 3: Variation of SINR across l with 2TX-2RX antennas and self reflector placed on the perpendicular axis, symmetrically between the two

reflected and leakage signal is odd multiple of π . Thus for $p \in \mathbb{Z}$,

$$\theta - 2\pi \frac{k}{\lambda} = (2p + 1)\pi$$

Substituting we get,

$$\theta\lambda - (2p + 1)\pi\lambda = 4\pi\sqrt{\frac{d^2}{4} + l^2}$$

Substituting $d = \lambda$, we get

$$\theta^2 + (4p^2 + 4p - 3)\pi^2 - (4p + 2)\pi\theta = \left(\frac{4\pi l}{\lambda}\right)^2$$

$$l = \frac{\lambda}{4\pi} \sqrt{(\theta - (2p + 3)\pi)(\theta - (2p - 1)\pi)}$$

The above expression is corroborated by the empirically observed periodic behavior of the troughs and peaks of cancellation.

4.2 Optimizing Self-Reflector Position

To maximize cancellation, one would want to identify the smallest distance l to maximize the cancellation power. Since $0 < \theta < 2\pi$, we see that l can exist either when both terms in the square root are positive or negative. Solving above, we get that solution always exists for $|p| > 1.5$. Yet, for a given θ , the smallest l can be obtained as follows

$$\begin{aligned} 0 < \theta < \pi : l &= \frac{\lambda}{4\pi} \sqrt{(\theta - 5\pi)(\theta - \pi)} & p = 1 \\ \pi < \theta < 2\pi : l &= \frac{\lambda}{4\pi} \sqrt{(\theta - \pi)(\theta + 3\pi)} & p = -1 \end{aligned}$$

Note how the above solution bounds the smallest $l < \frac{\sqrt{5}\lambda}{4}$ and the period between subsequent cancellation points decreases gradually. Further, as $\theta \rightarrow \pi$, the value of $l \rightarrow 0$ where it is infeasible to place the reflector. Even for this case, the next cancellation point l is bounded by $l < \frac{\sqrt{12}\lambda}{4}$. The above analysis highlights the three key benefits of our approach: (1) There will always be a solution, $l < 9.6\text{mm}$ for 27 GHz (the lowest ISM mm-Wave band), perfect for small-form devices. (2) We can achieve upto 10 dB SINR gain using our approach. (3) Continues to cancel the strong first tap of self-interference, even in multipath-rich environments.

Unknown θ : Measuring θ accurately in digital and feeding back to the movement platform remains too time-consuming for a real-world system. Another approach would be to brute-force search across the theoretical limits for a choice of l . Instead, we perform a detailed empirical analysis to show that the interference power is locally convex across the distance l (see Fig. 2(c)). This means we can perform gradient descent with hill climbing across first few lengths l to significantly speed up the process lowering latency for our system. Alg. 1 summarizes our approach.

Algorithm 1: Optimizing cancellation via self-reflector

Minimization function: $\Phi = \|Ae^{i\theta} + \frac{1}{k}e^{i2\pi\frac{k}{\lambda}}\|_2^2$

Output: l_{opt} optimum location of the self-reflector

```

1  $l = 0.5\lambda, \eta = 0.5$ 
2 while  $l > l_{min} \ \& \ l < \lambda \ \& \ \eta > 1e-3$  do
    | Gradient descent with  $l_{new} = l + \eta\Delta\Phi$ 
    |  $\eta_{new} = 0.75\eta$ 
3  $l_1 = l$ 
   //Hill Climbing if stuck in local minima
4 if  $l == \lambda$  then
    |  $l = 0$ , Repeat 2,  $l_2 = l$ 
5 if  $l == l_{min}$  then
    |  $l = \lambda$ , Repeat 2,  $l_2 = l$ 
6 if  $l_2 == NULL \mid \Phi(l_1) < \Phi(l_2)$  then
    |  $l_{opt} = l_1$ ;
   else
    |  $l_{opt} = l_2$ ;

```

Calibration, Sensitivity and effect on the desired SNR: Note that environmental changes only quite close (few cm) to the full-duplex node affect the location of our self-reflector since signal reflected from farther away objects are negligible in power compared to the dominant path of the self-interference signal leaking through the patch antenna substrate. Thus, the base station can periodically realign the self-reflector to compensate for these rare changes in the close proximity of the full duplex node. Further, this location is agnostic to the location of the transmitter at the other end. The green curve in Fig. 2(c) shows that the desired signal power does not vary much in the presence of the self-reflector. Fig. 2(c) and Fig. 3 also show that the self-reflector position is sensitive to 1-2 mm for optimal cancellation since the 3 dB peak width is around that value.

Generalization to phased arrays: Recent trends in mm-Wave are driven towards customized phased arrays that demonstrate high gains and directionality to improve the range and throughput of mm-Wave systems. Phased arrays add the complexity of not only removing interference between an antenna pair but the whole phased array ensemble put together. Thus, it would seem reasonable to argue that our approach does not extend to phased arrays.

However, remember our analysis in Fig. 2(b) demonstrated that the leakage decreases drastically as the distance between the antennas increases. Thus, even between the two phased arrays, where the antenna spacing between antennas is typically $\frac{\lambda}{2}$, the interference channel between the phased arrays will be dominated by the the

Figure 4: Simplified functional schematic of our chip

closest antennas which reduce to the above problem. We place 2 transmit and 2 receive antennas with self reflector placed symmetrically on the axis perpendicular to the line joining the 4 antennas and passing through the midpoint. Our empirical evaluation for the above described 2TX-2RX setup (shown in Fig.3) demonstrates remarkable similarity to the 1 transmit and 1 receive case showing feasibility of our approach to extend to phased arrays. Upon further inspection, we realize that this phenomenon is due to dominance of interference caused by the closest TX antenna for each RX antenna over all other TX antennas. Thus, this might allow our self-reflector to scale to other multi-antenna settings such as MIMO and phased arrays by developing a complex self-reflector along the symmetric axes between the transmit and receive antennas; although this remains to be verified experimentally. There have also been proposals for more complex phased array designs for full-duplex systems that intersperse TX-RX antennas and place them in 3D planes to exploit directionality, analysis of the correct choice of antennas' placement and design of self-reflector for such settings is left for future work.

5 MMFD ANALOG CANCELLATION

In this section, we describe how our custom chip provides enhanced mitigation of self-interference at the receive chain of the full-duplex system. Note that we need to significantly reduce the interference signal in analog to allow the analog-to-digital converter to capture the desired signal, which would otherwise be overpowered by one's own transmission.

5.1 Full-Duplex Analog Cancellation

Our first contribution to further cancel the self-interference is to develop one of the first full-duplex mm-Wave communication modules. Our work builds on some of the very first hybrid beamforming custom chips [52] that allow multi-antenna two-stream communication with efficient RF domain signal processing by using a fully connected structure. We propose an elegant approach to extend this fully connected hybrid beamforming transceiver architecture for full-duplex operation where one out of two streams can be re-purposed to inject the self-interference cancellation signal individually to every antenna while receiving through the other stream. This technique enables self-interference cancellation in a multi-antenna system without any hardware overhead. The significant advantage that our RF canceller provides over existing full-duplex cancellation techniques in mm-wave domain is that our structure

can be seamlessly scaled to multiple streams which is a norm in phased array systems.

A simplified functional schematic of the custom RF cancellation chip is shown in Fig. 4. The residual self-interference signal first passes through an impedance-matched low noise amplifier (LNA) to receive signal from the antenna with minimal reflection. The LNA amplifies the received signal minimizing degradation of the received SNR due to the circuit noise. The received signal after the LNA can be represented as:

$$y[\tau] = Ae^{i\theta} x_{tx}[\tau]$$

where x_{tx} is the transmitted signal, A is the coupling attenuation, and θ is the phase shift. Note, we do not assume a narrowband channel but instead aim to cancel the most dominant tap remaining after the antenna cancellation. Remember that, based on high attenuation of high frequency signals across distances, only the top dominant taps are responsible for most of the self-interference.

Our chip uses a weighted phase shifted version of the transmitted signal to cancel the interference in received signal after antenna cancellation. The RF canceller takes a copy of the transmitted signal as an input and applies a coarse-grain complex-weight, a fine-grain gain control coefficient and a fine-grain phase control angle to estimate the dominant tap of the self-interference channel. The coarse complex-weight is realized by using vector modulator; where a coupled-resonator based quadrature hybrid is used to generate quadrature signals. These quadrature signals are then independently weighted using a phase-invariant 5-bit (including 1-bit for sign control) variable gain amplifier pair followed by a combiner. The cancellation signal after applying coarse weights can be represented as:

$$c[\tau] = (W_{CI} + jW_{CQ})x_{tx}[\tau]$$

where the coarse complex-weights can achieve full 360° phase control. Our custom chip also features fine gain and phase adjustments that fine tune the above coarse adjustments to further improve the cancellation. The fine-gain control G_F is performed using a 5-bit variable gain amplifier whose gain can be linearly varied in a 30 dB range. The fine-phase control θ_F is performed by tuning a 5-bit capacitor bank of a resonator in the cancellation path. The fine phase control provides sub- 1° phase control with a full range of $\sim 20^\circ$. Therefore, the final cancellation signal can be represented as:

$$c[\tau] = (W_{CI} + jW_{CQ})G_F e^{j\theta_F} x_{tx}[\tau]$$

5.2 Designing Cancellation Weights

An important component of performing this analog cancellation is to choose the above canceller weights in a way that the interference signal component of the dominant tap is minimized. To achieve the above and automate the whole process, our chip provides a serial-to-parallel interface (SPI) which can control the aforementioned weights in the interference path. One would think that we can simply perform a brute force to estimate the above weights. However, an exhaustive search over the 20 bits of coarse and fine grained control leads to 2^{20} combinations which is particularly infeasible for low-latency communication required for mm-Wave. The optimization function minimizing the amount of interference power in the dominant tap (A_0) with received (Y_0) and transmitted (X_0)

signals can be written as:

$$\begin{aligned} \underset{\mathbf{W}}{\text{minimize}} \quad & \Omega_0 = \|A_0\|_2^2 \\ \text{where} \quad & A_0 = Y_0 - X_0(W_{CI} + jW_{CQ})G_F e^{j\theta_F} \\ & \mathbf{W} = \{W_{CI}, W_{CQ}, G_F, \theta_F\} \end{aligned}$$

In this work we have used a systematic search to significantly reduce the search space. In each searching cycle, we measure the average interference power after cancellation. We then use this measured power to feedback to our search algorithm. Note that while our current implementation uses a USRP X310 as the digital frontend, our approach generalizes to on-chip digital components. To reduce the search space, our algorithm relies on the redundancy provided by the fine-grained bits for accurately estimating the channel of the dominant tap. This allows us to instead sequentially choose each parameter independently. Based on the change in interference power, the algorithm checks the settings again to minimize interference.

Our algorithm is detailed in Alg. 2. Our optimization process reduces complexity from 2^{20} cycles to 81 cycles while achieving very close to the optimal performance. Our solution enables 18 dB mean interference cancellation across experiments in Sec. 9.1.

Algorithm 2: Optimizing cancellation in analog

Minimization function: $\Omega_0 = \|A_0\|_2^2$

Output: $\mathbf{W}_{opt}$ optimum weights

```

1  $W_{CI} = 1, W_{CQ} = 1, G_F = 15, \theta_F = 15$ 
2 for  $\theta = 0^\circ$  to  $360^\circ$  every  $30^\circ$  do
    | Update  $W_{CI}, W_{CQ}$  and measure  $\Omega_0$ 
    |  $W_{CI}, W_{CQ} = \underset{W_{CI}, W_{CQ}}{\text{argmin}} \Omega_0$  //12 cycles
3 for  $G = 1$  to  $31$  every  $2$  do
    |  $W_{CI}, W_{CQ} = \underset{W_{CI}, W_{CQ}}{\text{argmin}} \Omega_0$  //16 cycles
4 for  $\theta = \theta_{curr} - 30^\circ$  to  $\theta_{curr} + 30^\circ$  every  $3^\circ$  do
    |  $W_{CI}, W_{CQ} = \underset{W_{CI}, W_{CQ}}{\text{argmin}} \Omega_0$  //21 cycles
5 for  $G_F = G_{F,curr} - 5$  to  $G_{F,curr} + 5$  every  $1$  do
    |  $G_F = \underset{G_F}{\text{argmin}} \Omega_0$  //11 cycles
6 for  $\theta_F = \theta_{F,curr} - 10$  to  $\theta_{F,curr} + 10$  every  $1$  do
    |  $\theta_F = \underset{\theta_F}{\text{argmin}} \Omega_0$  //21 cycles

```

Another key challenge to make the above design work is to rapidly update the weight based on feedback from previous receptions. The faster the weight adaptation, the lower the latency of communication. Our solution relies on a key observation – the weight of the linear component changes in a specific pattern for static deployments. This pattern can be estimated by measuring the residual interference component in the digital domain. We use a feedback loop to continuously update the weight to mitigate change in weights due to linear effects of the interference over time.

6 MMFD DIGITAL CANCELLATION

In this section, we describe our approach to maximize the throughput by digital compensation of wireless impairments in the desired signal as well as removing residual interference from one's transmit chain. Our digital cancellation complements the analog cancellation via a continuous feedback loop described before and enables better SIC. Further, our digital cancellation can also attenuate multi-tap interference channel if required.

Previous full-duplex solutions in WiFi context for OFDM signals have primarily focused on removing the non-linear residual effects of the transmit signals in either the time or the frequency domain. These solutions have proven to be exceptionally effective at small bandwidths. Yet, at wider bandwidths, these techniques remain suboptimal or insufficient for removing interference effectively. Indeed, recent work in the ultra-wideband full-duplex domain[13] has shown large variance in phase and amplitude of the interference signal received by the full-duplex node.

A key reason why the above techniques don't work at the wider bandwidths (~GHz) of mm-wave frequencies is the large effects of even small timing offsets between transmit and received signals. Indeed, even on our hardware limited to 100MHz, the phase of the interference signal may vary by more than 2.5π across the bandwidth. This means conventional one-tap equalization and cancellation techniques while removing the interference at the center frequency may even amplify the interference at wider bands. Second, while prior work [13] attempts to perform frequency domain equalization in hardware effectively, digital domain provides significantly more resources to cancel the interference.

6.1 Dealing with Timing Offsets

Our approach to deal with timing offsets makes three observations about the interference signal, assuming that interferer and receiver operate using the same clock (a reasonable assumption as they are on the same full-duplex node): (1) The phase changes linearly across frequency at the same sample, dependent on the timing offsets; (2) The phase at a given frequency changes linearly across time, albeit slowly, due to tiny frequency offsets amplified during up-conversion; and (3) There are time-invariant noise distortions that are function of the transmitted signal that appear.

First, we demonstrate how we can compensate for the above hardware offsets even at low interference powers. The key idea behind our approach relies on the sparsity of phase variation information. Mathematically, the phase of the interference signal at any given frequency and time can be characterized as:

$$\phi_{f,t} = \phi_{0,0} + \alpha f + \beta t$$

where $\phi_{0,0}$ is the phase at center frequency at the start of the packet, α and β capturing the linear variation of phase with frequency and time respectively. Every symbol length, we get K subcarriers of phase information from the OFDM symbols. Further, for each symbol length, time increases by T, the symbol period of OFDM under consideration. Thus, for a long packet, lasting a time τ , one can use the phases to identify α, β and $\phi_{0,0}$. The above approach is resilient to frequency-selective fading as well as phase noise (Fig.5a).

Figure 5: mmFD Digital Cancellation: (a) Signal is chunked across frequency and time to estimate partial phases. The phase varies linearly across time and frequency. (b) Received signal after analog cancellation. (c) Signal after removing linear offsets. (d) Signal after removing time-invariant distortions

Algorithm 3: Removing linear offsets

Input: Received Signal y , Transmitted signal x

Output: Cancelled Received Signal z

```

1 forall numTaps do
2   forall  $t, f$  do
        $W_{f,t} = y_{f,t} / x_{f,t}$  //Measuring phase in chunks
        $\phi_{f,t} = \angle W_{f,t}$ 
        $A_{f,t} = |W_{f,t}|$ 
       //Using measured phase to fit linear model
3    $[\phi_{0,0}, \alpha, \beta, A_f] = \text{LinearModel}(\phi_{f,t}, A_{f,t})$ 
       //Cancelling based on linear model
4   forall  $t, f$  do
        $z_{f,t} = y_{f,t} - x_{f,t} A_f e^{i(\phi_{0,0} + \alpha f + \beta t)}$ 
5 Repeat until all taps are removed

```

Next, one would ideally want to estimate the amplitude of different symbols. One would think that amplitude would remain the same across frequency and time since it is typically not affected by timing and frequency offsets. However, this is not the case. First, the filters on software-defined radios, skewed antenna gain and unequal cancellation across frequencies means that amplitude is highly dynamic across frequencies. Second, in multipath rich environments, frequency selectivity could create large variations across frequencies. Thus, it is necessary for us to correctly identify these effects across frequency. Thus, the amplitude of the interference component can be characterized as

$$A_f = \sum_{i=0}^{\tau_f} A_{f,i}$$

Finally, how do we know where the interference signal lies in the received packet? Luckily, for us, the full duplex node completely knows the transmitted signal and can correlate with it to find the start of the packet. This can be further sped up by first correlating with the pilot symbols to identify the start of the packet. Alg. 3 details the approach described above to remove the linear offsets listed above (Fig.5b→5c).

6.2 Resolving Time-Invariant Distortions

Despite efficiently removing the linearly varying offsets, we are left with large amount of distortions created by various components. Indeed, [3] details how these distortions are heavily accentuated at such high frequencies due to components such as power amplifier, LNA, etc. These components saturate causing increased noise based on the transmitted signal power. Unfortunately, this relationship between the distortions and transmitted signal is not very deterministic.

We attempt to identify this relationship between the distortions and input signal. Across various signal powers and symbols transmitted, we notice two phenomena: (1) When transmitting the same signal with increasing powers, the distortions are highly correlated – i.e., learning these distortions at one power level would suffice to characterize them. (2) These distortions are time-invariant but change with transmitted data modulation, i.e., they are highly dependent on the data transmitted but not on when it is transmitted. We surmise that the source of these distortions in OFDM signals is the large peak-to-average power ratio variance across time. The above observations mean that with sufficient known transmitted data, one can reasonably estimate the behavior of these distortions inside an interference packet. Indeed, that is exactly what our digital frontend relies on to identify and eliminate the time-invariant distortions. Remember that requiring a downconverter with SDR to continuously monitor this link would cause an increase in cost and power consumption of each node.

Learning the time invariant distortion model: We first attempt to model the time-invariant distortions as a function across frequency and transmitted power. Our aim is to learn the mapping $Q : f, P \rightarrow k$ from frequency and transmitted symbol power to quantity of distortion as described below:

$$k = Q(f, P)$$

Since the mapping Q remains constant, we can learn this model by transmitting at various powers, OFDM symbols, in absence of the desired signal. Using Alg. 3, we will use the residual signal to learn the effects of time-invariant distortions on the symbols at that frequency. This model (a look-up table) can be verified and tested for each node during the manufacturing process and provided to purchaser along with the chip (or over the cloud).

Removing the time invariant distortions: We then emulate the effect that these distortions would have on the actual transmitted signal to create a distorted residual signal Γ . We then repeat Alg. 3 to correlate and remove this residual signal. Alg. 4 below describes our algorithm to cancel these distortions. (Fig.5c→5d)

Algorithm 4: Removing residual distortions

Input: Cancelled Signal z , Transmitted signal x ,
Residue mapping Q

Output: Digitally Cancelled Signal O

//Create distorted version of x

```

1 forall  $t, f$  do
  |  $\Gamma_{f,t} = Q(f, P(x(t)))$ 
  //Cancel distortion
2 Apply Alg. 3 with received signal as  $z$  and transmitted
  signal as  $\Gamma$ , with output as  $O$ 

```

Note that many components also generate non-linearities across time along with the background noise. Our digital cancellation techniques currently do not deal with the above impairments as they are much more complicated to decipher and estimate.

7 DISCUSSIONS AND LIMITATIONS

While our discussion thus far focuses on single antennas, this section describes how our system can be extended to phased arrays. We also describe the limitations of our system.

MIMO and Phased Arrays: As discussed in Sec. 4, our self-reflector design can be scaled to MIMO systems over phased arrays. However, this does not solve the problem completely since the residual interference would be a complex combination of interference from all the TX elements, and will still need to be cancelled at each of the RX streams. In order to push this combined self-interference to the noise floor, this would mean a quadratic increase in the processing complexity at both the analog and digital cancellation stages. However, as demonstrated in the [8], innovative methods using correlated self-interference channels and cascaded cancellation architecture can be explored for interference cancellation in MIMO to reduce the complexity. Our current chip's modular design [52] enables scaling to multi-antenna streams with RF SIC for each RX stream. While linear scaling will still be high for phased arrays with hundreds of elements, the analysis is beyond the scope of this paper since there are other effects which degrade the MIMO system in such large structures like spatial correlation between MIMO channels. Further, it has been shown in the WiFi context [65] that while gains of full duplex do not scale with MIMO configuration, full duplex remains useful to reduce the number of required streams of MIMO (ideally to half) for the same throughput performance.

Hardware Limitations: Our current implementation is bottlenecked by the limits of our off-the-shelf hardware to evaluate our system: (1) Our off-the-shelf software radios are limited to 100 MHz provided by USRP X310. (2) High distortion noise at mm-wave frequency adds an additional 10 dB noise figure to our setup[52]. (3) The maximum transmit power is limited by the low saturation

output power of the amplifiers (12 dBm) and the TRX board (0 dBm) that we use in the transmit path of the full-duplex node.

Scalability to higher bandwidths: The evaluation of our system is limited by the bandwidth of the SDR as mentioned above, we do believe that it can be extended to higher bandwidths. To verify this, we test the chip using a signal generator and the spectrum analyzer, allowing transmission and reception of higher bandwidth passband signal. We see that that chip is capable of achieving up to 26 dB average analog cancellation over a bandwidth of 500 MHz from 27.1 GHz to 27.6 GHz. While we cannot evaluate the performance of our digital canceller over a large bandwidth, theoretically, digital cancellation can support higher bandwidths due to the nature of its parallel design which is based on time and frequency chunking to deal with wideband signals. High bandwidth evaluation of end to end system is possible using multiple expensive signal generators and oscilloscopes, however, we have access to only a limited number of these expensive components and current state of the art SDRs are limited by their small bandwidths.

8 IMPLEMENTATION

Digital and SDR frontend: The full duplex node has two dedicated USRP X310s transmitting and receiving at 100 Msps sampling rate for transmit and receive chains, respectively. At its best, the USRP X310 is designed to operate at a maximum frequency of 6 GHz supporting a bandwidth of 120 MHz but our empirical tests showed that the X310s were able to reliably support only 100 MHz during transmission and reception without any packet drops as mentioned in the previous section. These USRPs are connected to 64-bit desktop computer running Ubuntu 18.04 OS via CAT 7 SFTP shielded ethernet cable designed for 10 Gigabit ethernet. To aid high speed data transfer between the computer and the USRP, a 10 Gigabit compatible 10Gtek X540-10G-T2-X8 ethernet card is installed in PCI Xpress3 port of the computer and a 10Gtek SFP adapter is installed in the RJ45 port of the USRP X310. We perform all our digital processing in MATLAB. For the PHY layer design, we use QAM modulated OFDM symbols with convolutional encoding at different coding rates. We develop custom Arduino code for updating the weights in our in-house chip[52] using SPI after learning them in MATLAB.

RF and Antenna frontend: To reach the desired frequency of 27.35 GHz, we design a custom in-house bidirectional TRX board (Fig. 7b) that upconverts the 1 GHz output from USRP to 27.35 GHz and vice versa. We use 2 RF-Lambda quadratic phase hybrids to interface the USRP with the TRX board, which generate separate I-Q streams from the an input complex baseband signal and vice-versa. In the transmit path, the high frequency signal is fed to the transmit patch antenna as well as a copy is fed to the the full-duplex chip for self-interference cancellation. In the receive path, the received signal (dominated by self-interference) from the receive patch antenna goes to the full-duplex chip for analog cancellation. Post analog cancellation, the signal is downconverted to baseband using the TRX board. We use **separate** Agilent 8251 20 GHz signal generators for generating the clock signal at each full-duplex node. For the self-reflector, we use linear motion stage with a resolution of 10 micrometers. The complete set-up at one of the full-duplex mmFD node is shown in Fig. 6a.

Figure 6: mmFD's implementation: (a) mmFD full duplex node, (b) Cancellation at various stages in presence of self-reflector for one instance of experiment, (c) Cancellation at various stages in absence of self-reflector for one instance of experiment

Figure 7: mmFD's hardware: (a) mmFD full duplex board, (b) TRX board for upconversion from 1 GHz to 27.35 GHz for transmit path and downconversion from 27.35 GHz to 1 GHz in the receive path

Testbed: We implement mmFD in an indoor lab setting of $10 \times 6 m^2$ consisting of various furniture and lab equipment which provide a rich multipath environment for the long range experiments. For the bidirectional experiments, we set up the two nodes on a bench top of length 1.5 meters surrounded by various metallic, plastic and wooden reflectors.

9 EVALUATION AND RESULTS

In this section, we evaluate our system across various parameters for unidirectional and bidirectional full-duplex communication.

9.1 Cancellation at various stages

Setup: We fix the transmit and receive gains of our antenna to maximum power without saturation. We then evaluate the ability of our system to cancel self-interference at various stages: (1) No cancellation; (2) With self-reflector; (3) With analog cancellation; (4) With digital cancellation. We use a 100 MHz OFDM modulated 64 QAM signal as the transmit signal. First, the self-reflector is optimized on the linear motion stage, followed by the analog and digital techniques. We measure the power of the noise signal across 100 MHz and compare it to the noise floor at -80 dBm (due to 10 dB noise figure of the hardware). The results are shown in Fig. 6b.

Antenna Isolation: We first notice that antenna isolation provides 30 dB interference suppression with the TX and RX antenna separated by one wavelength (1 cm) compared to the transmit power (black curve). TX-RX antenna separation reduces coupling

Component	Interference Cancellation
Antenna Isolation	30 dB
Self-Reflector	10 dB
Analog Cancellation	18 dB
Digital Cancellation	26 dB
Total Cancellation	84 dB

Table 1: Average SIC achieved at different stages

attenuating the self-interference power leaking into the RX path as shown in Fig. 2b in Sec.4.

Self-Reflector: Upon adding the self-reflector in front of the patch antenna and estimating optimum position using Alg. 1, we can achieve a further 9.88 dB cancellation. This shows the feasibility of using self-reflectors for interference cancellation in mm-Wave full-duplex systems

Analog Cancellation: Upon optimizing the weights using Alg. 2, we see a drop in self-interference power of 17.33 dB due to the cumulative effect of the self-reflector and analog cancellation showing ability of one to assist other.

Digital Cancellation: Upon adding digital cancellation to the chain, we see that the received interference power drops down by 21.35 dB for cumulative effect of 78.58 dB from the transmit power. Yet, the ability of our digital cancellation was indeed stifled because of large cancellations in the prior components. We see in Fig. 6c that once one removes the self-canceller, the digital canceller can still reduce the interference power by 30.37 dB. Thus the system is capable of achieving similar performance in the presence and absence of the self-reflector. This is not to depict that the self-reflector is not useful, but instead to show that when a higher transmit power is used, we can add the self reflector in Fig. 6c and still reach the noise floor after cancellation. Table 1 summarizes the average cancellation capabilities of each component across multiple experiments. Note that Fig. 6b and 6c are just representative of two such instances of such experiments.

Extension to longer ranges: Our primary evaluation for SIC at 0 dBm was shown in Fig. 6b and 6c. However, to test the range limits of mmFD, we evaluate it's SIC in the presence of higher transmit powers using a high saturation power amplifier (HMC1132PM5E) in the transmit path. We increase the average transmit power for the

Figure 8: mmFD's unidirectional evaluation: (a) Average SNR difference between cancelled signal and received signal without interference in various setups. (b) Throughput of mmFD vs. the ideal and half-duplex throughput. (c) Throughput across distance in LOS, BLOCK and NLOS setups compared to signals without interference

Figure 9: SINR vs. Transmit Power of mmFD

Figure 10: mmFD unidirectional evaluation setup

Figure 11: BER for unidirectional setups vs. QAM

mmFD node from -6 dBm to 21 dBm in steps of 3 dB. We keep the desired SNR at 20 dB and measured the SINR after cancellation. We see in Fig. 9 that although SINR decreases marginally till 12 dBm due to the non-linearities of the power amplifier, there is a sudden decrease by 3 dB at 15 dBm transmit power which we believe is the limit of our SIC.

9.2 Unidirectional full-duplex link

In this section, we evaluate the ability of mmFD's cancellation in presence of desired transmission in a multipath-rich environment.

Setup: We use a transmitter with a horn antenna placed at a known distance away from a mmFD node. We use Vector Telekom horn

antenna to overcome the range limitations and access greater flexibility of directionality. We transmit different combinations of 16, 32, and 64 QAM OFDM symbols with convolutional encoding with rates $\frac{1}{2}$, $\frac{2}{3}$ and $\frac{3}{4}$ at 100 Msps from the USRP. For throughput calculations, we choose the best QAM and coding rate combination to maximize throughput while ensuring BER less than 10^{-4} . Based on attenuation across distance, our desired signal's antenna's transmit signal is at most 10 dB more powerful than the transmit(self interference) signal of the full duplex mmFD node. We consider 2 specific cases of non-line-of-sight (NLOS) setting to compare our system's performance with direct line of sight case. In the first NLOS case called 'BLOCK', we block Line of Sight path between the horn antenna transmitter and the FD node using a 1 cm thick wooden plank with the horn antenna still pointing at the mmFD node. In the second NLOS case called 'NLOS', we point the horn antenna transmitter towards a steel plank which acts as a reflector. The mmFD node then receives the reflected signal bouncing off the reflector. The 3 setups can be seen in Fig. 10.

Performance across distance: For this experiment, we move the horn antenna from 2 meters to 8 meters in steps of 1 meter from the fixed full duplex node inside a laboratory setting. Fig. 8b shows the throughput across distance for the mmFD node. While our cancellation techniques reduce the interference power drastically, the effect of the remaining interference prevents us from achieving ideal performance. However, note that we still do far superior compared to the case when the one way link was used as TDD or FDD for a half-duplex communication.

LOS vs. NLOS: A key objective of our system relies on ensuring that the SINR of the received signal after cancellation reaches as close to that of the signal without any self-interference. Fig. 8a shows the difference in SNR across distances. The small average error in SNR compared to transmitted signal shows the efficacy of our self-interference cancellation.

Throughput: We next evaluate the throughput of our client for distances from 2 to 8 m (Fig. 8c). As expected, we observe a steady decrease in bitrate as the distance grows larger. Note that baseline appears to perform better than mmFD as it represents optimal half duplex bitrates in absence of any self interference compared to mmFD unidirectional links in the presence of self interference. Perhaps more surprising is the fact that the BLOCK experiments

Figure 12: mmFD bidirectional evaluation: (a) Average SINR difference between cancelled and normal received signal, (b) Throughput of mmFD vs. half-duplex communication, (c) Bit error rate across coding rates for 64 QAM

provide higher throughput than NLOS experiments despite the antenna steering to avoid the blockage. This highlights the heavy attenuation that mm-Wave signals face as they propagate across distances, so much that the attenuation of a wooden board is preferable to a meter longer path.

Bit Error Rates: Finally, we assess the bit error rate across modulation schemes (Fig. 11) for the 3 scenarios. As expected the errors increase with distance and the BER is worse in case of NLOS compared to BLOC across all modulation schemes. This trend follows that of the throughput showing drastic increases in errors across modulation schemes.

9.3 Bidirectional full-duplex link

In this section, we evaluate our system in the first real-world deployment of bidirectional full duplex link.

Setup: For the bidirectional full duplex Link, we installed two mmFD nodes on a laboratory benchtop of length 1.5 m. We transmit different combinations of 16, 32, and 64 QAM OFDM symbols with convolutional encoding with rates $\frac{1}{2}$, $\frac{2}{3}$ and $\frac{3}{4}$ at 100 Msps from the USRP. We then find the Bit Error Rates (BER) at the receiver node after digital processing. We then choose the QAM and coding rate combination that yields the best data reception rate whose BER is less than 10^{-4} . For the half duplex node, we do not transmit from the node which is receiving at a given time instant, whereas for the full duplex node, we transmit and receive at both the nodes simultaneously. In order to find the throughput of the full duplex node, we add the individual throughput of the two FD nodes of the link. We then increase the separation of the nodes from 50 cm to 110 cm in steps of 10 cm. The low maximum transmit power (0 dBm) due to saturation issues described in Sec. 7 as well as the low gain of the patch antennas limit us from going beyond this range. However, our evaluation in Fig. 9 shows the capability of our system to work at longer ranges if a higher transmit power is used by effectively cancelling the additional self-interference.

SNR Difference: Fig. 12a shows the average difference of full-duplex node SNR (after SIC) from that of half-duplex equivalent at the three distances. The average difference is around 0.1 dB and remains consistently low across distance. The low residual noise power shows the efficacy of our self-interference cancellation upto the noise level.

Throughput: Fig. 12b shows that our full-duplex system achieves significantly larger throughput over conventional half-duplex systems. Remember that despite having a higher throughput one-way, both node will either have to perform FDD or TDD to share resources. This will lead to effective halving of the one-way throughput. Our results show that mmFD can achieve upto $1.7\times$ the one-way throughput and 752 Mbps mean 2-way communication throughput.

Bit Error Rates: Fig. 12c shows the average BER achieved for the two nodes at 3 different distances for 64 QAM at 1/2, 2/3 and 3/4 coding rates. We see that although the full duplex nodes at the two sides of the link are supposed to be identical in performance, however, that may not necessarily be true: it should be noted that the channel from node 1 to node 2 is not symmetric. In fact, these are two different channels altogether. One might feel that the TX and RX antennas being very close together might lead to symmetric channel, yet this is not true at mm-wave frequencies where even a sub-degree(angle) change in antenna orientation can drastically change the channel.

10 CONCLUSION

This paper presents, mmFD, the first comprehensive evaluation of bidirectional full-duplex link at mm-wave frequencies. mmFD shows challenges posed in enabling full-duplex communication at 28 GHz via SDR-driven study. Based on the observed wireless impairments, mmFD presents novel solutions in antenna design, analog cancellation and digital cancellation which specifically exploit opportunities offered by mm-Wave through its small form-factor, high directionality and heavy attenuation. Our results demonstrate SIC of 84 dB achieving upto $1.7\times$ the half-duplex throughput. While our work focuses on a single link PHY-layer design, we believe mmFD opens the door for future work in mm-Wave full-duplex that studies its impact on the MAC and higher-layer network and transport protocols as well as phased array and MIMO networks.

Acknowledgements: We thank the shepherd and reviewers for their insightful feedback which helped improve this paper. We also thank all the members of the WiTech Lab at CMU for their valuable inputs. We are grateful to Professor Bob Iannucci from CMU Silicon Valley campus for promptly providing hardware support for our experiments. This research was supported in parts by NSF grant 1823235, Kavčić-Moura grant and IoT@CyLab.

REFERENCES

- [1] ABBAS, H., AND HAMDI, K. Full duplex relay in millimeter wave backhaul links. In *2016 IEEE Wireless Communications and Networking Conference* (2016), IEEE, pp. 1–6.
- [2] ABDOLLAHI, S., AL-RAWESHIDY, H., FAKHRAIE, S. M., NILAVALAN, R., AND KAMAREI, M. Full duplex 60 GHz millimeter wave transmission over multi-mode fiber. In *2010 Second International Conference on Ubiquitous and Future Networks (ICUFN)* (2010), IEEE, pp. 222–226.
- [3] AHMED, E., AND ELTAWIL, A. M. All-digital self-interference cancellation technique for full-duplex systems. *IEEE Transactions on Wireless Communications* 14, 7 (2015), 3519–3532.
- [4] ARIYARATHNA, V., MADANAYAKE, A., TANG, X., COELHO, D., CINTRA, R. J., BELOSTOTSKI, L., MANDAL, S., AND RAPPAPORT, T. S. Analog approximate-FFT 8/16-beam algorithms, architectures and CMOS circuits for 5G beamforming MIMO transceivers. *IEEE Journal on Emerging and Selected Topics in Circuits and Systems* 8, 3 (2018), 466–479.
- [5] ARRIBAS, E., ANTA, A. F., KOWALSKI, D., MANCUSO, V., MOSTEIRO, M., WIDMER, J., AND WONG, P. W. Optimizing mmwave wireless backhaul scheduling. *IEEE Transactions on Mobile Computing* (2019).
- [6] BEAS, J., CASTANON, G., ALDAYA, I., ARAGÓN-ZAVALA, A., AND CAMPUZANO, G. Millimeter-wave frequency radio over fiber systems: a survey. *IEEE Communications surveys & tutorials* 15, 4 (2013), 1593–1619.
- [7] BHARADIA, D., JOSHI, K., AND KATTI, S. Robust full duplex radio link. In *ACM SIGCOMM Computer Communication Review* (2014), vol. 44, ACM, pp. 147–148.
- [8] BHARADIA, D., AND KATTI, S. Full Duplex MIMO Radios. In *11th USENIX Symposium on Networked Systems Design and Implementation (NSDI 14)* (2014), pp. 359–372.
- [9] BHARADIA, D., McMILIN, E., AND KATTI, S. Full duplex radios. In *ACM SIGCOMM computer communication review* (2013), vol. 43, ACM, pp. 375–386.
- [10] CERWALL, P., JONSSON, P., MÖLLER, R., BÄVERTOFT, S., CARSON, S., AND GODOR, I. Ericsson mobility report. *On the Pulse of the Networked Society*. Hg. v. Ericsson (2015).
- [11] CHAFFEE, D. L., MALLORY, M. P., AND BRAND, G. H. Frequency domain adaptive echo canceller for full-duplex data transmission, May 26 1992. US Patent 5,117,418.
- [12] CHEN, B., WU, Y., HAN, M., AND ZHANG, Q. A Novel Architecture of Millimeter-Wave Full-Duplex Radio-over-Fiber System with Source-Free BS Based on Polarization Division Multiplexing and Wavelength Division Multiplexing. *Progress In Electromagnetics Research* 80 (2018), 103–110.
- [13] CHEN, T., DASTJERDI, M. B., ZHOU, J., KRISHNASWAMY, H., AND ZUSSMAN, G. Wideband full-duplex wireless via frequency-domain equalization: Design and experimentation. *arXiv preprint arXiv:1812.01126* (2018).
- [14] CHI, T., PARK, J. S., LI, S., AND WANG, H. A 64GHz full-duplex transceiver front-end with an on-chip multifeed self-interference-canceling antenna and an all-passive canceler supporting 4Gb/s modulation in one antenna footprint. In *2018 IEEE International Solid-State Circuits Conference-(ISSCC)* (2018), IEEE, pp. 76–78.
- [15] CHOI, J. I., JAIN, M., SRINIVASAN, K., LEVIS, P., AND KATTI, S. Achieving single channel, full duplex wireless communication. In *Proceedings of the sixteenth annual international conference on Mobile computing and networking* (2010), ACM, pp. 1–12.
- [16] CUI, H., MA, M., SONG, L., AND JIAO, B. Relay selection for two-way full duplex relay networks with amplify-and-forward protocol. *IEEE Transactions on Wireless Communications* 13, 3 (2014), 3768–3777.
- [17] DE MELO GUIMARÃES, L., AND LUIZ BORDIM, J. A full-duplex mac tailored for 5g wireless networks. *Wireless Communications and Mobile Computing* 2018 (2018).
- [18] DINC, T., CHAKRABARTI, A., AND KRISHNASWAMY, H. A 60 GHz CMOS full-duplex transceiver and link with polarization-based antenna and RF cancellation. *IEEE journal of solid-state circuits* 51, 5 (2016), 1125–1140.
- [19] DUARTE, M., SABHARWAL, A., AGGARWAL, V., JANA, R., RAMAKRISHNAN, K., RICE, C. W., AND SHANKARANARAYANAN, N. Design and characterization of a full-duplex multiantenna system for WiFi networks. *IEEE Transactions on Vehicular Technology* 63, 3 (2013), 1160–1177.
- [20] EL AYACH, O., RAJAGOPAL, S., ABU-SURRA, S., PI, Z., AND HEATH, R. W. Spatially sparse precoding in millimeter wave MIMO systems. *IEEE transactions on wireless communications* 13, 3 (2014), 1499–1513.
- [21] ESHRAGHI, A., BORDES, J. P., DAVIS, C., RAO, R. K., ALI, M., AND DENT, P. W. Vehicular radar system with self-interference cancellation, Oct. 17 2017. US Patent 9,791,551.
- [22] FAN, L., LI, Y., AND ZHAO, M. Joint iq imbalance and pa nonlinearity pre-distortion for highly integrated millimeter-wave transmitters. In *2014 IEEE Globecom Workshops (GC Wkshps)* (2014), IEEE, pp. 399–404.
- [23] FEBRIANTO, T., HOU, J., AND SHIKH-BAHAEL, M. Cooperative full-duplex physical and mac layer design in asynchronous cognitive networks. *Wireless Communications and Mobile Computing* 2017 (2017).
- [24] GHASEMPOUR, Y., DA SILVA, C. R., CORDEIRO, C., AND KNIGHTLY, E. W. IEEE 802.11 ay: Next-generation 60 GHz communication for 100 Gb/s Wi-Fi. *IEEE Communications Magazine* 55, 12 (2017), 186–192.
- [25] GHOSH, S., AND SEN, D. An Inclusive Survey on Array Antenna Design for Millimeter-wave Communications. *IEEE Access* (2019).
- [26] HAMAZUMI, H., IMAMURA, K., LAI, N., SHIBUYA, K., AND SASAKI, M. A study of a loop interference canceller for the relay stations in an SFN for digital terrestrial broadcasting. In *Globecom'00-IEEE. Global Telecommunications Conference. Conference Record (Cat. No. 00CH37137)* (2000), vol. 1, IEEE, pp. 167–171.
- [27] HAN, S., ZHANG, Y., MENG, W., AND ZHANG, Z. Precoding design for full-duplex transmission in millimeter wave relay backhaul. *Mobile Networks and Applications* 23, 5 (2018), 1416–1426.
- [28] HASSANIEH, H., ABARI, O., RODRIGUEZ, M., ABDELGHANY, M., KATABI, D., AND INDYK, P. Fast millimeter wave beam alignment. In *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication* (2018), ACM, pp. 432–445.
- [29] HEATH, R. W., GONZALEZ-PRELCIC, N., RANGAN, S., ROH, W., AND SAYEED, A. M. An overview of signal processing techniques for millimeter wave MIMO systems. *IEEE journal of selected topics in signal processing* 10, 3 (2016), 436–453.
- [30] HUSSAIN, M. T., SHARAWI, M. S., PODILCHACK, S., AND ANTAR, Y. M. Closely packed millimeter-wave MIMO antenna arrays with dielectric resonator elements. In *2016 10th European Conference on Antennas and Propagation (EuCAP)* (2016), IEEE, pp. 1–4.
- [31] JAIN, M., CHOI, J. I., KIM, T., BHARADIA, D., SETH, S., SRINIVASAN, K., LEVIS, P., KATTI, S., AND SINHA, P. Practical, real-time, full duplex wireless. In *Proceedings of the 17th annual international conference on Mobile computing and networking* (2011), ACM, pp. 301–312.
- [32] JIA, Z., YU, J., AND CHANG, G.-K. A full-duplex radio-over-fiber system based on optical carrier suppression and reuse. *IEEE Photonics Technology Letters* 18, 16 (2006), 1726–1728.
- [33] JOG, S., WANG, J., GUAN, J., MOON, T., HASSANIEH, H., AND CHOUDHURY, R. R. Many-to-many beam alignment in millimeter wave networks. In *16th USENIX Symposium on Networked Systems Design and Implementation (NSDI 19)* (2019), pp. 783–800.
- [34] KIM, C.-Y., KIM, J.-G., AND HONG, S. A quadrature radar topology with Tx leakage canceller for 24-GHz radar applications. *IEEE transactions on microwave theory and techniques* 55, 7 (2007), 1438–1444.
- [35] KIM, D., LEE, H., AND HONG, D. A survey of in-band full-duplex transmission: From the perspective of phy and mac layers. *IEEE Communications Surveys & Tutorials* 17, 4 (2015), 2017–2046.
- [36] KUMARI, P., CHOI, J., GONZÁLEZ-PRELCIC, N., AND HEATH, R. W. IEEE 802.11 ad-based radar: An approach to joint vehicular communication-radar system. *IEEE Transactions on Vehicular Technology* 67, 4 (2017), 3012–3027.
- [37] KUTTY, S., AND SEN, D. Beamforming for millimeter wave communications: An inclusive survey. *IEEE Communications Surveys & Tutorials* 18, 2 (2015), 949–973.
- [38] LI, B., ZHAO, C., SUN, M., ZHANG, H., ZHOU, Z., AND NALLANATHAN, A. A bayesian approach for nonlinear equalization and signal detection in millimeter-wave communications. *IEEE Transactions on Wireless Communications* 14, 7 (2015), 3794–3809.
- [39] LI, L., JOSIAM, K., AND TAORI, R. Feasibility study on full-duplex wireless millimeter-wave systems. In *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (2014), IEEE, pp. 2769–2773.
- [40] LI, L., REYNAERT, P., AND STEYAERT, M. S. Design and analysis of a 90 nm mm-wave oscillator using inductive-division lc tank. *IEEE Journal of Solid-State Circuits* 44, 7 (2009), 1950–1958.
- [41] LIAO, Y., BIAN, K., SONG, L., AND HAN, Z. Full-duplex mac protocol design and analysis. *IEEE Communications Letters* 19, 7 (2015), 1185–1188.
- [42] LIN, K., WANG, Y. E., PAO, C.-K., AND SHIH, Y.-C. A $K\alpha$ -band FMCW radar front-end with adaptive leakage cancellation. *IEEE Transactions on Microwave Theory and Techniques* 54, 12 (2006), 4041–4048.
- [43] MA, J. 5 Gbit/s full-duplex radio-over-fiber link with optical millimeter-wave generation by quadrupling the frequency of the electrical RF carrier. *IEEE/OSA Journal of Optical Communications and Networking* 3, 2 (2011), 127–133.
- [44] MA, BOJIANG AND SHAH-MANSOURI, HAMED AND WONG, VINCENT WS. Full-duplex relaying for D2D communication in millimeter wave-based 5G networks. *IEEE Transactions on Wireless Communications* 17, 7 (2018), 4417–4431.
- [45] MARALI, D., AND GURBUZ, O. Design and performance analysis of a full-duplex mac protocol for wireless local area networks. *Ad Hoc Networks* 67 (2017), 53–67.
- [46] MARZI, Z., RAMASAMY, D., AND MADHOW, U. Compressive channel estimation and tracking for large arrays in mm-wave picocells. *IEEE Journal of Selected Topics in Signal Processing* 10, 3 (2016), 514–527.
- [47] MO, J., SCHNITZER, P., PRELCIC, N. G., AND HEATH, R. W. Channel estimation in millimeter wave MIMO systems with one-bit quantization. In *2014 48th Asilomar Conference on Signals, Systems and Computers* (2014), IEEE, pp. 957–961.
- [48] MONDAL, S., AND PARAMESH, J. A reconfigurable 28-/37-ghz mmse-adaptive hybrid-beamforming receiver for carrier aggregation and multi-standard mimo communication. *IEEE Journal of Solid-State Circuits* 54, 5 (2019), 1391–1406.
- [49] MONDAL, S., SINGH, R., HUSSEIN, A. I., AND PARAMESH, J. A 25-30 ghz 8-antenna 2-stream hybrid beamforming receiver for mimo communication. In *2017 IEEE Radio Frequency Integrated Circuits Symposium (RFIC)* (2017), IEEE, pp. 112–115.
- [50] MONDAL, S., SINGH, R., HUSSEIN, A. I., AND PARAMESH, J. A 25–30 ghz fully-connected hybrid beamforming receiver for mimo communication. *IEEE Journal*

- of *Solid-State Circuits* 53, 5 (2018), 1275–1287.
- [51] MONDAL, S., SINGH, R., AND PARAMESH, J. A reconfigurable 28/37ghz hybrid-beamforming mimo receiver with inter-band carrier aggregation and rf-domain lms weight adaptation. In *2018 IEEE International Solid-State Circuits Conference-(ISSCC)* (2018), IEEE, pp. 72–74.
 - [52] MONDAL, S., SINGH, R., AND PARAMESH, J. 21.3 a reconfigurable bidirectional 28/37/39ghz front-end supporting mimo-tdd, carrier aggregation tdd and fdd/full-duplex with self-interference cancellation in digital and fully connected hybrid beamformers. In *2019 IEEE International Solid-State Circuits Conference-(ISSCC)* (2019), IEEE, pp. 348–350.
 - [53] MOON, T., GAUN, J., AND HASSANIEH, H. Online Millimeter Wave Phased Array Calibration Based on Channel Estimation. In *IEEE VLSI Testing Symposium (VTS)* (2019), IEEE.
 - [54] MUMTAZ, S., RODRIGUEZ, J., AND DAI, L. *mmWave Massive MIMO: A Paradigm for 5G*. Elsevier Science, 2016.
 - [55] PAN, H. K. Mm-wave phased array antenna with beam tilting radiation pattern, Aug. 22 2017. US Patent 9,742,077.
 - [56] PARK, J., HONG, J.-P., AND BEAK, S. Optimal beamforming with limited feedback for millimeter-wave in-band full-duplex mobile X-haul network. *IEEE Access* 6 (2018), 51038–51048.
 - [57] PERAHIA, E., CORDEIRO, C., PARK, M., AND YANG, L. L. IEEE 802.11 ad: Defining the next generation multi-Gbps Wi-Fi. In *2010 7th IEEE consumer communications and networking conference* (2010), IEEE, pp. 1–5.
 - [58] PIKSA, P., ZVANOVEC, S., AND CERNY, P. Elliptic and hyperbolic dielectric lens antennas in mmwaves. *Radioengineering* 20, 1 (2011), 270–275.
 - [59] RAJAGOPAL, S., TAORI, R., AND ABU-SURRA, S. Self-interference mitigation for in-band mmwave wireless backhaul. In *2014 IEEE 11th Consumer Communications and Networking Conference (CCNC)* (2014), IEEE, pp. 551–556.
 - [60] RASEKH, M. E., MARZI, Z., ZHU, Y., MADHOW, U., AND ZHENG, H. Noncoherent mmWave path tracking. In *Proceedings of the 18th International Workshop on Mobile Computing Systems and Applications* (2017), ACM, pp. 13–18.
 - [61] RIIHONEN, T., BALAKRISHNAN, A., HANEDA, K., WYNE, S., WERNER, S., AND WICHMAN, R. Optimal eigenbeamforming for suppressing self-interference in full-duplex MIMO relays. In *CISS* (2011), pp. 1–6.
 - [62] SANGIAMWONG, J., ASAI, T., HAGIWARA, J., OKUMURA, Y., AND OHYA, T. Joint multi-filter design for full-duplex MU-MIMO relaying. In *VTC Spring 2009-IEEE 69th Vehicular Technology Conference* (2009), IEEE, pp. 1–5.
 - [63] SATYANARAYANA, K., EL-HAJJAR, M., KUO, P.-H., MOURAD, A., AND HANZO, L. Hybrid beamforming design for full-duplex millimeter wave communication. *IEEE Transactions on Vehicular Technology* 68, 2 (2018), 1394–1404.
 - [64] SHIRAZI, A. H. M., NIKPAIK, A., MOLAVI, R., LIGHTBODY, S., DJAHANSHAHI, H., TAGHIVAND, M., MIRABBASI, S., AND SHEKHAR, S. On the design of mm-wave self-mixing-vco architecture for high tuning-range and low phase noise. *IEEE Journal of Solid-State Circuits* 51, 5 (2016), 1210–1222.
 - [65] SULTAN, R., SONG, L., SEDDIK, K. G., AND HAN, Z. Full-duplex meets multiuser mimo: Comparisons and analysis. *IEEE Transactions on Vehicular Technology* 66, 1 (2016), 455–467.
 - [66] SUN, S., AND RAPPAPORT, T. S. Millimeter wave mimo channel estimation based on adaptive compressed sensing. In *2017 IEEE International Conference on Communications Workshops (ICC Workshops)* (2017), IEEE, pp. 47–53.
 - [67] SUN, S., RAPPAPORT, T. S., AND SHAFT, M. Hybrid beamforming for 5G millimeter-wave multi-cell networks. In *IEEE INFOCOM 2018-IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)* (2018), IEEE, pp. 589–596.
 - [68] SUR, S., ZHANG, X., RAMANATHAN, P., AND CHANDRA, R. BeamSpy: enabling robust 60 GHz links under blockage. In *13th USENIX Symposium on Networked Systems Design and Implementation (NSDI 16)* (2016), pp. 193–206.
 - [69] WANG, P., PAJOVIC, M., ORLIK, P. V., KOIKE-AKINO, T., KIM, K. J., AND FANG, J. Sparse channel estimation in millimeter wave communications: Exploiting joint aod-aoa angular spread. In *2017 IEEE International Conference on Communications (ICC)* (2017), IEEE, pp. 1–6.
 - [70] WEI, T., ZHOU, A., AND ZHANG, X. Facilitating robust 60 ghz network deployment by sensing ambient reflectors. In *14th USENIX Symposium on Networked Systems Design and Implementation (NSDI 17)* (2017), pp. 213–226.
 - [71] XIAO, Z., XIA, P., AND XIA, X.-G. Full-duplex millimeter-wave communication. *IEEE Wireless Communications* 24, 6 (2017), 136–143.
 - [72] ZHOU, A., ZHANG, X., AND MA, H. Beam-forecast: Facilitating mobile 60 GHz networks via model-driven beam steering. In *IEEE INFOCOM 2017-IEEE Conference on Computer Communications* (2017), IEEE, pp. 1–9.
 - [73] ZHOU, J., CHUANG, T.-H., DINC, T., AND KRISHNASWAMY, H. 19.1 receiver with >20MHz bandwidth self-interference cancellation suitable for fdd, co-existence and full-duplex applications. In *2015 IEEE International Solid-State Circuits Conference-(ISSCC) Digest of Technical Papers* (2015), IEEE, pp. 1–3.
 - [74] ZHU, Y., ZHU, Y., ZHAO, B. Y., AND ZHENG, H. Reusing 60ghz radios for mobile radar imaging. In *Proceedings of the 21st Annual International Conference on Mobile Computing and Networking* (2015), ACM, pp. 103–116.