

Wideband Real-Time Adaptive Full-Duplex Radios

Alon S. Levin¹, Igor Kadota², Bo Zhang¹, Sasank Garikapati¹,
Manav Kohli¹, Harish Krishnaswamy¹, Mingoo Seok¹, Gil Zussman¹

¹Electrical Engineering, Columbia University, ²Electrical Engineering, Northwestern University
{alon.s.levin, bz2358, sg3734, mpk2138, hk2532, ms4415, gz2136}@columbia.edu, kadota@northwestern.edu

Abstract

Full-duplex (FD) wireless enables simultaneous transmission and reception on the same frequency channel, enhancing spectrum and power efficiency provided self-interference (SI) is reliably suppressed. Despite significant research into SI cancellation (SIC), real-world FD deployment remains an open challenge. The highly dynamic nature of SI, especially at higher bandwidths, requires system-level designs capable of adapting to a changing electromagnetic environment in order to sustain cancellation. We present a wideband, adaptive FD radio that combines a highly configurable RFIC canceller with a novel, device-specific, sparsity-based configuration algorithm for real-time SIC tracking. The proposed data-driven algorithm exploits measured tap responses of the RFIC canceller to account for hardware imperfections and is implemented on an FPGA as part of an SDR-based FD radio. In VNA-based validation, this algorithm outperforms theory-based and literature-inspired baselines across bandwidths up to 100 MHz. In the SDR-based FD radio, the FPGA-controlled RFIC canceller with linear digital SIC achieves 72 dB of over-all cancellation across 100 MHz, reducing SI to the noise floor. Under mobility, dynamic optimization outperforms static routines and maintains noise-floor-level cancellation despite variations in the SI channel. Finally, an FD link between two SDRs exhibits <0.8 dB EVM degradation relative to an equivalent half-duplex link, confirming that a bidirectional FD link can approximately double throughput.

CCS Concepts

• **Hardware** → **Wireless devices**; • **Networks** → **Physical links**; **Network algorithms**; **Network experimentation**.

Keywords

Full-duplex wireless, software-defined radios, self-interference cancellation, adaptive optimization

ACM Reference Format:

Alon S. Levin, Igor Kadota, Bo Zhang, Sasank Garikapati, Manav Kohli, Harish Krishnaswamy, Mingoo Seok, and Gil Zussman. 2026. Wideband Real-Time Adaptive Full-Duplex Radios. In *The 32nd Annual International Conference on Mobile Computing and Networking (MobiCom '26)*, October 26–30, 2026, Austin, TX, USA. ACM, New York, NY, USA, 19 pages. <https://doi.org/10.1145/3795866.3844155>

1 Introduction

Full-duplex (FD) radios simultaneously transmit and receive on the same frequency, unlike half-duplex (HD) radios, which separate these two operations with frequency-division or time-division duplexing. FD wireless can, therefore, improve spectral efficiency at the physical (PHY) layer as well as provide benefits to higher network layers [41, 54, 64]. Achieving successful FD operation is challenging due to the extremely strong self-interference (SI) signal present whenever the transmitter (Tx) is active. Approximately 70–110 dB of SI cancellation (SIC) is required at the receiver (Rx) to achieve noise floor-level cancellation and preserve the signal-to-interference-plus-noise ratio (SINR); typically, SIC spans the antenna interface, analog, and digital domains [8, 13, 40, 49].

An FD radio's SI channel is highly dynamic, particularly during mobility as the SI multipath coupling structure is continuously reshaped [40]. This necessitates an analog cancellation stage capable of tracking and adapting to the evolving SI channel, for any misconfiguration of the analog stage risks degrading the communication link or, in more severe cases, saturating the Rx's frontend entirely, at which point digital SIC can no longer recover the desired incoming signal. It is critical for a real-time adaptive analog cancellation stage to keep pace with the SI channel's evolution under mobility.

This challenge grows with channel bandwidth, especially when considering the need for small-form-factor RF integrated circuit (RFIC)-suitable analog SI cancellers for real-world mobile device deployment. Analog cancellation performance degrades sharply with increasing bandwidth, requiring complex RFIC architectures with more independent parameters that can accurately model the SI channel's multipath structure across the wider band [26, 49]. Precisely configuring these parameters in real time is non-trivial: the configuration space grows exponentially, and tuning must account for both the instantaneous SI channel and device-specific hardware imperfections [47]. Despite meaningful




This work is licensed under a Creative Commons Attribution 4.0 International License.

MobiCom '26, Austin, TX, USA


© 2026 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2505-0/2026/10

<https://doi.org/10.1145/3795866.3844155>



(a) Wideband real-time adaptive FD radio.



(b) Microphotograph of the RFIC canceller.

Figure 1: Implemented FD radio and its canceller.

progress in IC design [13], RFIC capability alone is insufficient: system-level challenges surrounding practical deployment have thus far prevented widespread FD adoption [56]. Computationally efficient algorithms capable of adaptively configuring and maintaining cancellation as the channel evolves without disrupting the data link are therefore a prerequisite for practical FD deployment.

We address these challenges by integrating a highly configurable wideband TDE RFIC canceller [48] — previously shown to achieve strong SIC over 80 MHz under idealized, non-communicating lab conditions [26] — as the frontend of a Software-Defined Radio (SDR)-based FD radio (Figure 1). To enable real-time adaptive RF SIC, we present a novel globally- and locally-sparse reconstruction-based configuration algorithm, implemented on an FPGA parallel to the FD node’s signal processing workflow. We evaluate the performance of the algorithm and the node in realistic scenarios, including mobility and an FD link. *To the best of our knowledge, this is the first real-world implementation and experimental evaluation of a wideband FD radio design configured by a sparsity-employing algorithm, and the first optimization and evaluation of SIC under mobile conditions with a real FD radio.*

Organization. The remainder of this work is organized as follows. Section 2 reviews related work. Section 3 briefly covers the RFIC canceller’s operating principles, background used to formulate the configuration problem and two solutions in Section 4 which are then validated against literature baselines in Section 5. Section 6 describes the FD radio implementation and experimentation, with discussion in Section 7. Section 8 follows with concluding remarks.

2 Related Work

FD wireless concepts (e.g., canceller designs, radio implementations, and performance gains) have been extensively surveyed in [13, 39, 49, 54], including its potential for beyond-5G/6G networks [30, 56]. Below, we summarize relevant works, and present a few key FD radio designs in Table 1.

RFIC Cancellor Design. RF SIC is achieved by (i) maximizing SI isolation at the antenna interface and (ii) employing dedicated RF SIC circuits to actively cancel SI before the signal reaches the analog-to-digital converter (ADC) at the Rx, preventing saturation and desired signal desensitization [49].

Tx-Rx isolation has been pursued via dual-antenna (bistatic) schemes [5, 15, 19, 31, 52, 60] and shared-antenna interfaces such as circulators [8, 17, 25, 33, 53] or electrically balanced duplexers [20, 42, 46]. Massive-antenna systems can utilize beamforming to achieve Tx-Rx isolation [4, 24].

Numerous RF SI cancellation techniques have been proposed, including (i) amplitude and phase-based (P&A) cancellers [14, 22, 55, 65], (ii) time-domain equalization (TDE) cancellers that operate by equalizing the delay of the received SI [6, 8, 16, 63], and (iii) frequency-domain equalization (FDE) cancellers that operate by equalizing the spectral components of the SI channel [12, 27, 66]. P&A cancellers are effectively degenerate TDE cancellers, with frequency-flat amplitude- and phase-scaling that can only emulate the SI channel at a single frequency point [13]. In contrast, TDE cancellers enable wideband cancellation by employing a bank of delay line elements that emulate the frequency-selective multipath SI channel, at the cost of increased configuration complexity as the number of tunable taps grows. Wideband, RFIC-suitable implementations of TDE and FDE cancellers have been enabled by switched-capacitor delay lines [48] and N-path filters [34], respectively.

System-Level FD Radio Design. Prior work has evaluated FD SI cancellers acting as the frontend of an SDR [8, 14, 35, 55], including network-level FD testbeds [36]. System-level designs that employ an adaptive RF SI canceller commonly implement the control algorithm on a host PC [2, 12, 37] or an FPGA [38, 57]. RF canceller designs with few parallel paths may be configured using least-mean-squares [2, 8, 12], but wideband cancellers with many paths crucial to wideband operation require more computationally efficient methods, such as simplified circuit-based models [66] or polynomial fitting of measured frequency responses [26, 48]. Several prior works simulate mobility by moving items about a static PCB canceller connected to laboratory test equipment [2, 40, 57]. Such scenarios, though time-varying, feature the perturbation of relatively few individual scatterers within the larger, static multipath environment. A mobile node, however, experiences a fully changing landscape in which all components of the SI channel evolve jointly — a structurally different

Table 1: System-Level Comparison of Prior Full-Duplex Radio Designs

Work	Real-Time Analog Control & Mobility						Experiment Setup & Results		
	Class / HW	Taps	Tunable Params	Compute / RF Platform	Compute Time	Mobile Eval.	Bandwidth / Center Freq.	Avg. Tx. Power	Iso. + Analog / Total SIC
[8]	TDE PCB	16	Gain only	Software Test Equip.	$\sim 920 \mu\text{s}$	No	80 MHz 2.45 GHz	20 dBm	62 dB 110 dB
[14]	P&A PCB	1	Gain+Phase	Software SDR	N.S.	No	5 MHz 0.9 GHz	0 dBm	40 dB 90 dB
[38]	TDE PCB	4	Gain+Phase	Software Test Equip.	$\sim 500 \mu\text{s}$	No	120 MHz 2.45 GHz	30 dBm	70 dB 70 dB*
[40]	TDE PCB	3	Gain+Phase	N.S. Test Equip.	N.S.	Simul.	80 MHz 2.46 GHz	8 dBm	63 dB 88 dB
[12]	FDE PCB	2	Gain+Phase +Filter	Software SDR	$< 10 \text{ ms}$	No	20 MHz 0.9 GHz	10 dBm	52 dB 95 dB
[2]	P&A PCB	2	Gain+Phase	Software Test Equip.	N.S.	Simul.	10 MHz 2.442 GHz	N.S.	35 dB* 35 dB*
[55]	P&A RFIC	1	Gain+Phase	Offline SDR	N. S.	No	100 MHz 27.35 GHz	0 dBm	58 dB 84 dB
[37]	P&A PCB	2	Gain+Phase	Software Test Equip.	N.S.	No	20 MHz 2.5 GHz	N.S.	48 dB* 48 dB*
[57]	P&A PCB	3	Gain+Phase	FPGA Test Equip.	N.S.	Simul.	80 MHz 2.46 GHz	8 dBm	63 dB 63 dB*
This Work	TDE RFIC	16	Gain+Delay	FPGA SDR	Initial: $\sim 94 \text{ ms}$ Adjust: $\sim 735 \mu\text{s}$	Yes	100 MHz 1.05 GHz	-10 dBm	45 dB 72 dB

N.S.: Not Stated in source paper. *Partial sum: isolation and/or digital stages not stated or included.

and more complex problem. This work instead combines a custom RFIC canceller, real-time FPGA control, and an SDR into a single deployable radio, evaluated under true mobility.

This paper extends our previously published demos [43, 44], showcasing early prototype results, with a detailed treatment of the algorithm and hardware implementation, along with extensive evaluation across multiple experiments. Furthermore, it extends previously-reported capabilities in order to demonstrate higher bandwidths and mobile operation.

3 Background: The RFIC Canceller

In this section, we briefly describe the operating principles of the RFIC canceller developed in [48], shown in Figure 1b. For additional details, we refer the reader to Appendix A. The RFIC canceller is integrated into the frontend of the FD radio on a signal path parallel to that of the antenna interface, and acts as a filter aiming to match the SI channel experienced at the antenna interface. The signal received at the Rx antenna, consisting of the desired signal and the multipath SI signal, is combined with the Tx signal filtered through the RFIC; in an ideal outcome, the RFIC canceller perfectly inverts the SI channel, resulting in complete destructive interference of the SI signal and fully isolating the desired signal.

The RFIC canceller is implemented in a standard 64 nm CMOS process. Tx power is capacitively coupled into the RFIC canceller, which is composed of 16 switched-capacitor-based delay taps. The taps fall into three classes as shown in

Table 2: (i) one zero-delay tap with direct input-output connection, (ii) five low-delay and (iii) ten high-delay taps, providing up to 2 ns and 8 ns of delay respectively at 250 ps resolution over a 1 GHz bandwidth. The tap outputs are weighted, summed, and injected into the Rx to achieve cancellation. Intuitively, the RFIC canceller realizes an adaptive FIR filter with sixteen parallel taps of tunable gain and delay.

4 Adaptation Algorithm

The large number of possible gain and delay configurations renders a manual configuration scheme infeasible. The time-varying environment within which the RFIC canceller must adapt demands an efficient algorithm that runs in real-time. The challenge of determining a low-complexity algorithmic solution is compounded by the manufacturing process variation between individual chips. We provide a basic formulation of this problem in Section 4.1, followed by two proposed solutions: a theory-based approach in Section 4.2 and a data-driven approach in Section 4.3. For additional details, we refer the reader to Appendix B.

4.1 Basic Problem Formulation

The RFIC canceller consists of 16 taps, each having a programmable gain α_n selected from a discrete set of values normalized between 0 and 1, and a programmable delay τ_n selected from a discrete set of values $\{\tau_n^{(0)}, \dots, \tau_n^{(P_n)}\}$, where

Table 2: RFIC Canceller Taps

Tap Class	Count	Parameter	Settings	Maximum
Zero-Delay	1	Gain	64 (6 bits)	-27 dB
		Delay	1 (0 bits)	—
Low-Delay	5	Gain	64 (6 bits)	-33 dB
		Delay	8 (3 bits)	2 ns
High-Delay	10	Gain	64 (6 bits)	-36 dB
		Delay	32 (5 bits)	8 ns

P_n denotes the maximum delay index for tap n , determined by the class of the tap as shown in Table 2.

We denote the frequency response of the n^{th} RF tap as $H_n(f, \theta_n)$, where $\theta_n = (\alpha_n, \tau_n)$ represents the selected canceller parameters for the tap; the complete set of configuration parameters is given by $\theta = \{(\alpha_n, \tau_n)\}_{n=0}^{15} \in \Theta$, where Θ is the *admissible configuration space* containing every permissible parameter combination. We note that Θ has over 10^{19} delay-selection combinations alone.¹

Let $H_{\text{SI}}(f)$ and $H_{\text{RFIC}}(f, \theta)$ represent the frequency responses of the SI channel and the RFIC canceller, configured with parameters $\theta \in \Theta$, respectively. The goal of a configuration algorithm is to find the optimal configuration θ^* which maximizes SIC or, equivalently, minimizes the *residue*:

$$\begin{aligned} \text{Res}(f, \theta) &:= H_{\text{SI}}(f) - H_{\text{RFIC}}(f, \theta), \\ \theta^* &:= \arg \min_{\theta \in \Theta} \|\text{Res}(f, \theta)\|^2. \end{aligned} \quad (1)$$

4.2 Theory-Based Formulation & Solution

The RFIC taps were designed to achieve ideal frequency responses, with the n^{th} RF tap characterized by

$$H_n(f, \theta_n) = \alpha_n e^{-j2\pi f \tau_n}. \quad (2)$$

Assuming that the tap's true response closely matches this ideal analytical model, the RFIC canceller's response can be modelled as the sum of the individual tap responses, and the optimization problem in Equation 1 becomes

$$\theta^* := \arg \min_{\theta \in \Theta} \left\| H_{\text{SI}}(f) - \sum_{n=0}^{15} \alpha_n e^{-j2\pi f \tau_n} \right\|^2. \quad (3)$$

Projected Gradient Descent (PGD) can be used to solve this optimization, using the iterative update function

$$\theta_{t+1} \leftarrow \text{Proj}[\theta_t - \eta_t \nabla g(\theta_t)], \quad (4)$$

where η_t is the learning rate, $g(\theta)$ is the objective in Equation 1, and $\text{Proj}[\theta]$ is an operator that projects an out-of-bounds configuration θ onto the nearest in-bound limit point (i.e., $\alpha_n \in [0, 1]$). The PGD algorithm is performed iteratively until a stopping criterion on the gradient size is achieved. Appendix B includes a closed-form expression for the gradient $\nabla g(\theta_t)$ (Equation 18) and a pseudo-algorithm (Algorithm 2).

¹This can be shown by combining the number of delay options each tap has, based on its class (see Table 2). This calculation yields $1^1 \times 8^5 \times 32^{10} \approx 10^{19.6}$.




Figure 2: Ideal versus VNA-measured frequency responses of the 8th RF tap with delay $\tau_8^{(25)} = 25 \times 0.25$ ns. Polynomial models fail to match the true behavior.

Once a solution θ^* is found, the final step is to map the configuration of the ideal model to the real hardware. A naive mapping would program θ^* directly onto the RFIC canceller. A better approach is to map each *selected* ideal tap configuration $H_n(f, \theta_n^*)$ to the best-fit *measured* individual tap response $\hat{H}_n(f, \tau_n)$, obtained via an offline tap characterization campaign in which each tap $n \in \{0, \dots, 15\}$ is measured alone at each delay setting $\{\tau_n^{(0)}, \dots, \tau_n^{(P_n)}\}$. This mapping uses a least-squares approximation to select physical taps with strong fits to the ideal, after which the mapped configuration is programmed onto the RFIC canceller.

4.3 Data-Driven Formulation & Solution

A main drawback of the theory-based solution stems from the assumption that the frequency responses of the taps in a practical IC are identical to their ideal models. Figure 2 shows the deviations in tap response that can be sufficiently significant as to render the proposed solution inadequate for obtaining sufficient cancellation from the RFIC canceller.

Rather than using the ideal tap response defined in Equation 2, we can instead define a *data-driven* model in which the response of tap n with configuration θ_n is a linear scaling of the tap's *measured response* $\hat{H}_n(f, \tau_n)$ when configured with fixed unit gain and delay τ_n , as follows:

$$H_n(f, \theta_n) = \alpha_n \hat{H}_n(f, \tau_n). \quad (5)$$

The frequency response of the RFIC canceller across all sixteen RF taps is therefore modeled as a linear combination

$$H_{\text{RFIC}}(f, \theta) = \sum_{n=0}^{15} \alpha_n \hat{H}_n(f, \tau_n). \quad (6)$$

To ensure that the linear model (used in Equations 3 and 6) approximates the true performance of the RFIC canceller, we impose the following constraints: (i) the sum of the gains is below a threshold K , and (ii) the number of enabled taps is below a threshold M . The importance of these constraints to the algorithm's performance is later showcased in Table 3.

$$\sum_{n=0}^{15} \alpha_n \leq K \quad (7) \quad \sum_{n=0}^{15} \mathbb{1}\{\alpha_n > 0\} \leq M \quad (8)$$

To determine the optimal taps, delays, and gains, we propose a two-phase configuration algorithm, consisting of an *initial selection* stage followed by *iterative adjustments*.

Initial Selection. Given the measured frequency response of the SI channel $H_{\text{SI}}(f)$, our goal is to determine an “adequate” θ that reduces SI significantly while keeping the RFIC canceller’s operation in the linear region (i.e., subject to Equations 7 and 8). This formulation is effectively a sparse reconstruction problem, wherein we have a goal of representing the measured SI channel using a sparse linear combination of measured tap responses at fixed delays.

Consider the block matrix $\mathbf{H}$ composed of submatrices $\mathbf{H}_n$, each representing the *measured frequency responses* of tap n over its possible delay settings:

$$\mathbf{H} = [\mathbf{H}_0 \cdots \mathbf{H}_{15}], \text{ where } \mathbf{H}_n = [\hat{H}_n(f, \tau^{(p)})]_{p=0}^{P_n} \quad (9)$$

As each tap can only have a single delay setting at any given time, the gain of tap n can be represented by a sparse column vector $\mathbf{A}_n$ with at most one non-zero entry, equal to the selected gain α_n at index p , corresponding to delay $\tau_n^{(p)}$. This results in an additional per-tap sparsity constraint:

$$\mathbb{1}\{\mathbf{A}_n > \mathbf{0}\} \leq 1. \quad (10)$$

The individual tap gain vectors can be concatenated into a single sparse vector representing the gains of all RF taps:

$$\mathbf{A} = [\mathbf{A}_0^\top \cdots \mathbf{A}_{15}^\top]^\top. \quad (11)$$

The response of the RFIC canceller can thus be represented as the matrix product of the *dictionary of measured tap responses* $\mathbf{H}$ and the gain vector $\mathbf{A}$, i.e., $H_{\text{RFIC}}(f, \mathbf{A}) = \mathbf{H}\mathbf{A}$. Correspondingly, the optimization problem in Equation 1 can now be expressed as an explicit sparse reconstruction:

$$\begin{aligned} \mathbf{A}^* &:= \arg \min_{\mathbf{A}} \|\mathbf{H}_{\text{SI}}(f) - \mathbf{H}\mathbf{A}\|^2 \\ \text{s.t. } &\|\mathbf{A}\|_1 \leq K \\ &\|\mathbf{A}\|_0 \leq M \\ &\|\mathbf{A}_n\|_0 \leq 1, \forall n \in \{0, \dots, 15\} \end{aligned} \quad (12)$$

Note that there are three distinct sparsity constraints: (i) a *net sum* constraint (Equation 7), (ii) a *global sparsity* constraint (Equation 8), and (iii) a *local sparsity* constraint (Equation 10). These constraints enable the indirect selection of appropriate delay settings despite only returning a solution for the gain vector, as the support of the solution $\mathbf{A}^*$ corresponds directly to the selected (tap, delay) tuples in the optimal configuration θ^* .

Significant literature has been devoted to determining algorithmic approaches to sparse reconstruction problems [21], but as the problem presented in Equation 12 includes explicit conditions on the size and structure of the support, the family of greedy support-seeking algorithms emerges as a clear contender. We take inspiration from *Orthogonal Matching*

Pursuit (OMP), an efficient greedy algorithm that selects support elements one at a time, given a particular selection criterion, such as the maximal correlation between the unselected dictionary elements and the residual signal [58]. Though less powerful than convex-relaxation alternatives [10] and not guaranteed to recover the signal perfectly, greedy algorithms’ reduced computational complexity suits real-time adaptive analog SIC, especially given that it is sufficient to cancel enough to avoid Rx saturation.

We present *Globally-Locally Constrained OMP (GLC-OMP)* in Algorithm 1, a greedy sparsity-ensuring reconstruction algorithm that selects taps, delays, and gains given an arbitrary SI channel measurement. The iterative structure of the greedy algorithm allows us to address the global and local sparsity constraints explicitly. The number of steps is limited to the number of desired enabled taps M , as a single (tap, delay) tuple is added to the support of the solution at each step’s conclusion. Once a tap is selected, its remaining tuples are excluded from further consideration, ensuring that no tap has multiple delays selected for inclusion — an additional constraint that traditional sparse reconstruction algorithms do not take into consideration. The optimal (tap, delay) tuple at any given step in the optimization is that whose inclusion best minimizes the SI residue. Once M taps are selected, their respective (tap, delay) tuples and gains are mapped to the sparse gain vector $\mathbf{A}$, which is then quantized to the discrete values allowed by the RFIC canceller.

Algorithm 1 Globally-Locally Constrained OMP

Input: SI channel $H_{\text{SI}}(f)$

- 1: $\mathcal{S} \leftarrow \emptyset$ $\triangleright$ Selected tap-delay tuples.
- 2: $\mathcal{U} \leftarrow \{(n, p) \mid n \in \{0, \dots, 15\}, p \in \{0, \dots, P_n\}\}$ $\triangleright$ Unselected.
- 3: **while** $|\mathcal{S}| < M$ **do**
- 4: **for** $(i, j) \in \mathcal{U}$ **do**
- 5: $\mathcal{S}_{\text{aux}} \leftarrow \mathcal{S} \cup \{(i, j)\}$
- 6: Find $\{\alpha_n\}$ for all $(n, p) \in \mathcal{S}_{\text{aux}}$ that optimizes:
- $$\text{SI}(i, j) = \min_{\substack{0 \leq \{\alpha_n\} \leq 1 \\ \sum \alpha_n \leq K}} \left\| \hat{H}_{\text{SI}}(f) - \sum_{(n, p) \in \mathcal{S}_{\text{aux}}} \alpha_n \hat{H}_n(f, \tau_n^{(p)}) \right\|_2^2 \quad (\star)$$
- 7: **end for**
- 8: $(i^*, j^*) \leftarrow \arg \min \{\text{SI}(i, j)\}$ $\triangleright$ Find best-fit tap-delay tuple.
- 9: $\mathcal{S} \leftarrow \mathcal{S} \cup \{(i^*, j^*)\}$
- 10: $\mathcal{U} \leftarrow \mathcal{U} \setminus \{(i^*, p) \mid p \in \{0, \dots, P_{i^*}\}\}$ $\triangleright$ Exclude tap i^* .
- 11: **end while**
- 12: Solve $(\star)$ to obtain final gains $\{\alpha_n\}$ for all $(n, p) \in \mathcal{S}$.
- 13: Construct gain vector $\mathbf{A} = [\mathbf{A}_0^\top \cdots \mathbf{A}_{15}^\top]^\top$, where

$$[\mathbf{A}_n]_p = \begin{cases} \alpha_n, & \text{if } (n, p) \in \mathcal{S}. \\ 0, & \text{otherwise.} \end{cases}$$

14: **return** Quantize($\mathbf{A}$)

Iterative Adjustment. After an initial configuration is obtained from GLC-OMP, it is necessary to evaluate and improve on its performance for two reasons: (i) despite the constraints introduced to maintain operation in the linear region, there is no guarantee that the hardware will perform exactly as intended, potentially resulting in poorer-than-anticipated performance; and (ii) the SI channel is highly dynamic, requiring tracking and adaptivity to maintain steady SIC.

One option is to compute a new configuration using GLC-OMP for every measured SI channel. However, even taking into account the low complexity of greedy algorithms, this is highly inefficient as the algorithm requires checking every (tap, delay) tuple; furthermore, this still does not address the potential mismatch between the linearity model and true hardware performance. Instead, we can take inspiration from the theory-based approach and modify the PGD algorithm, as shown in Appendix B, to adjust the gain vector $\mathbf{A}$, keeping the support (i.e., the selected taps and delays) constant. Under this simplification, the gradient and update step become

$$\nabla g(\mathbf{A}) = -2 \Re \left[\mathbf{H}^H (H_{\text{SI}}(f) - \mathbf{H}\mathbf{A}) \right], \quad (13)$$

$$\mathbf{A}_{t+1} \leftarrow \text{Proj}[\mathbf{A}_t - \eta_t \nabla g(\mathbf{A}_t)]. \quad (14)$$

This algorithm can be adapted to find solutions to the constrained linear least squares in Algorithm 1; this is a desirable implementation as it removes the need for additional optimization tools and reduces computation time.

5 Algorithm Validation

In this section, we evaluate the performance of the algorithms of Section 4. Section 5.1 describes the validation testbed, as well as two algorithms from the literature used as baseline comparisons. Section 5.2 then presents the results of the experiments and discusses their implications.

5.1 VNA-Based Validation Setup

Figure 3 shows the block diagram of the testbed used to validate the previously discussed optimization algorithms. A VNA measures $H_{\text{SI}}(f)$ and communicates it to a PC running MATLAB. An optimization script, using any of the proposed or baseline approaches, computes the next canceller configuration, and the RFIC canceller is configured using a micro-controller acting as a serial interface adapter between the PC and the canceller. These incremental configuration adjustments are repeated for a fixed number of iterations, and the maximum RF cancellation achieved is recorded. The VNA is configured to measure 2001 data points, from which we select a bandwidth of interest. All storage and computation is performed in double-precision floating-point format (64 bits). This setup provides a well-controlled environment for isolating and analyzing the performance of online configuration algorithms in a repeatable and comparable manner.




Figure 3: Block diagram of the validation testbed.

Initial isolation results from a bistatic Tx/Rx interface [51] using two commercial off-the-shelf (COTS) swivel whip antennas, whose physical separation provides higher initial RF isolation across a wider bandwidth than possible with single-antenna interfaces [49]. The transmit antenna is connected directly to the transmit port of the RFIC canceller, and the receive antenna is connected to a power combiner so that the incoming SI signal destructively interferes with the output signal of the RFIC canceller.

To compare the proposed theory-based and data-driven approaches to literature baselines, we additionally implement the cubic polynomial modeling approach described in [26], and a delay-selection approach inspired by [8]. We briefly describe these methods here; for additional details, we refer the reader to Appendix B.

Cubic Polynomial Modeling. In [26], each measured (tap, delay) response is fit to a cubic polynomial model, as shown in Figure 2; these models are used in a greedy algorithm followed by gradient-descent refinement to optimize the RFIC canceller. Two important differences from our approach: (i) the cubic model for measured tap responses, and (ii) the greedy algorithm, which is not based on OMP.

Straddled-Delay Selection. In [8], the authors perform an intentional, hardware-driven delay selection: fixed delay lines are spaced to straddle the SI channel's principal delay, producing lagging and leading copies of the SI signal that combine linearly for cancellation. We emulate this by fixing delay settings for uniform spacing around the observed SI delay, then solving a linear least-squares problem for gains with iterative gradient-descent refinement.


5.2 Results and Analysis

We implemented all four algorithms to optimize the RFIC canceller over four bandwidths (10 MHz, 25 MHz, 50 MHz, 100 MHz) centered at 1050 MHz. The residual SI channel was recorded at every iteration, tracking overall RF SIC throughout the experiments. A sweep over the sum of gains threshold K and the number of enabled taps M was performed during the data-driven optimization.

Table 3 presents the SIC due to the RFIC canceller using all four algorithms, including the optimal hyperparameters for the data-driven approach and its performance when the

Table 3: Validation of Algorithmic Performance – Improvement in RF SIC

Bandwidth	Proposed Algorithms				Baseline Algorithms	
	Theory-Based	Data-Driven	(K^*, M^*)	$K = M = 16$	Cubic Fit	Delay-Select
10 MHz	25.4 dB	37.7 dB	(0.8, 4)	11.8 dB	25.9 dB	24.5 dB
25 MHz	16.1 dB	31.1 dB	(1.2, 6)	14.7 dB	14.8 dB	22.3 dB
50 MHz	10.7 dB	26.1 dB	(1.2, 6)	16.1 dB	12.7 dB	12.1 dB
100 MHz	5.8 dB	14.9 dB	(1.6, 4)	2.7 dB	6.9 dB	7.1 dB

**Figure 4: Residual SI channel at each step of the data-driven optimization over four bandwidths. In each experiment, the optimal SIC achieved (solid blue) was greater than that achieved using the theory-based approach (magenta).**

constraints are fully relaxed (i.e., $K = M = 16$). It can be seen that the proposed data-driven approach achieves significantly greater overall RF SIC than the benchmarks. Figure 4 shows the residual SI at each stage of the data-driven optimization approach with optimal hyperparameter selection for all four bandwidths, highlighting the channel after initial isolation, after the GLC-OMP stage of the optimization, and at the optimal RFIC canceller configuration; the optimal residual channel after theory-based optimization is overlaid.

The evolution of the overall RF SIC at each step of the proposed optimization processes is shown in Figure 5. Beyond poorer overall performance, the theory-based optimization shows worse convergent behavior, particularly at higher bandwidths. Although the 10 MHz experiment converged to a local minimum, the remaining three diverged from the optimal configuration, resulting in poorer SIC with clear oscillatory behavior in the 25 MHz and 50 MHz cases. By comparison, the data-driven optimization converged to the optimal configuration for all four bandwidths.

The optimal hyperparameters trend toward improved performance with fewer enabled taps M and a smaller sum of gains K , confirming the assumptions made when imposing

the optimization constraints (Equations 7 and 8). In particular, we see that if we fully relax the constraints on M and K , as would be assumed from a pure circuit theory model, the performance gains collapse. As GLC-OMP is greedy and each additional tap selection incurs additional runtime overhead, we reduce unnecessary computation by placing a hard upper bound on the number of enabled taps at $M = 6$ when deploying the algorithm in a practical implementation.

When evaluating the performance of baseline algorithms, it becomes clear that there is a significant advantage in making use of the exact measured performance of the RFIC canceller taps. Simplifying assumptions, such as ideal linear models or even high-order polynomial models, are insufficient to capture the non-idealities present in the system, resulting in performance loss. By the same token, by attempting to match the measured principal SI delay with average group delays of different (tap, delay) tuples, the delay-straddling approach loses information about the non-ideal and frequency-selective amplitude and phase response of individual taps, which is critical, especially as the bandwidth increases. Furthermore, the delay-straddling approach does not effectively translate to the case of a bistatic antenna




Figure 5: Evolution of the SIC at each iteration of the theory-based and data-driven optimizations.

interface, as multipath effects due to environmental factors negate the argument of a single primary delay path, as there would be through the circulator used in [8].

Finally, we note that the proposed data-driven approach aligns naturally with current industry practices in RFIC calibration and deployment [3]. The ATE-driven measurement mechanisms used to build tuning algorithms that account for silicon process variation are mirrored by the methodology presented in this work – namely, the algorithmic use of device-specific measurements to obtain performance gains over traditional circuit model-based approaches.

6 Implementation and Evaluation

In this section, we describe the implementation and evaluation of a wideband SDR-based FD radio with real-time adaptive RF SIC utilizing the data-driven algorithm described in Section 4.3. Section 6.1 discusses the architecture of the FD radio, including PHY layer considerations, the FPGA-based implementation of the optimization algorithm, and digital SIC. Experimentation with the adaptive FD radio and discussions of the results are then presented in the ensuing sections, with node-level evaluation of the system’s overall SIC capabilities presented in Section 6.2 for a stationary node and in Section 6.3 in mobility, and an evaluation of the FD radio’s performance in a wideband link in Section 6.4.

6.1 SDR-Based FD Radio

Figure 6 shows the block diagram of the adaptive FD radio. The VNA is replaced with a USRP X310 SDR, which streams complex samples to and from a host PC over a 10 GbE LAN interface. All data and control flow is performed through a GNU Radio flowgraph running UHD [23, 28], which implements a custom OFDM-based PHY layer for transmit/receive DSP functions and digital SIC [45]. The computation of the optimal canceller configuration, and by extension, the configuration process, is delegated to an FPGA running a custom hardware image. Details and insights on the implementation, including optimizations and design choices, are presented in Appendix C; code and documentation is available in [18].




Figure 6: Block diagram of the SDR-based FD radio.

OFDM-Based PHY Layer. We transmit WiFi-like packets with a Zadoff-Chu (ZC) pilot preamble for robust synchronization and SI channel estimation [29, 32]. The data is pre-generated, modulated, and packetized using an open-source GNU Radio module for the IEEE 802.11 standard [9]. We use 256-sample OFDM symbols with 242 non-zero subcarriers, and a cyclic prefix ratio of 0.25 (64 samples). In our experiments, the transmitted frame structure consisted of four pilot slots (two slots uniquely allotted to each radio in a link), followed by fifteen data slots that are utilized by both radios simultaneously. This yields a per-radio duty cycle of $17/19 \approx 89.5\%$, and an FD pilot overhead of $4/19 \approx 21.1\%$.²

ZC sequences belong to a class of constant amplitude, zero autocorrelation (CAZAC) waveforms, taking the form $b_k = \exp(j\pi k(k+1)r/N)$ with length N and root r [50]. These sequences are utilized in various capacities by the 3GPP LTE standard [1]. In particular, their zero out-of-phase autocorrelation enables the precise synchronization necessary to (i) accurately estimate the magnitude and phase characteristics of $H_{SI}(f)$ in consistent fashion, and (ii) properly synchronize the outgoing and incoming packet streams so that the data fields overlap, achieving true FD operation.

System Control. Real-time system control is accomplished through a custom C++ out-of-tree (OOT) module running in GNU Radio, designed to support FD operation [45]. As packets are received, the pilot symbols are extracted and fed into a linear estimator in order to obtain $H_{SI}(f)$. The SI channel estimates, with an appended operation instruction, are sent to the FPGA to compute the next configuration. Once the canceller is configured, a new channel estimate is obtained, and the control loop continues until convergence or a set number of iterations; a generalized system-level block diagram of this process is shown in Figure 7. Additional mechanisms were developed to account for hardware failures due to the high sampling rate requirements.

A key design consideration was the decoupling of the communication and optimization functions to avoid introducing unnecessary latency that would affect packet flow; receive-side SIC optimization should not hinder the packet transmission process. To achieve this, the system controller

²The number of data slots reflect an experimental design choice rather than an upper bound; increasing this number per frame reduces pilot overhead.




Figure 7: FPGA-based optimization loop, running in parallel to the GNU Radio flowgraph’s processes.

communicates with the FPGA in parallel with the Tx and Rx signal chains, asynchronously polling channel state information when the FPGA requires updated data for computation. **FPGA Computing and Control Hardware.** We designed custom blocks on a Zynq FPGA to compute the optimal configuration and configure the RFIC canceller using the data-driven algorithms presented in Section 4.3. The FPGA’s processing system (PS) consists of the UART port and ARM core, whereas the programming logic (PL) contains the computation and configuration hardware blocks.

The computation hardware consists of the ROM_{tap} block, which stores the 242-point tap measurements at every delay setting, and the ConfigOpt block, which uses the measurements to compute the next configuration. Specifically, the measurements are composed of the characterizations of each tap at every delay setting, and are obtained in a one-time per-device and per-bandwidth offline measurement campaign using the FD radio under standard operating procedures; environmental effects are removed in post-processing via all-taps-disabled baseline measurements, providing a clean performance dataset that we can use for online environmental adaptation. All data, including the stored tap measurements and SI channel estimates sent to the FPGA, are stored in Q1.15 fixed-point format (16 bits). The configuration hardware consists of a memory block, Mem_{cfg} , which buffers the 750-bit binary configuration stream, and a canceller configuration block, CanConfig , which configures the canceller chip over a serial scan chain.

The ARM core of the FPGA runs a control program that supports three operational modes: manual system configuration, GLC-OMP-based computation, and PGD-based computation. The GNU Radio flowgraph running on the PC triggers these operations by sending a command to the ARM core over UART, followed by any additional parameters necessary for the operation to run, such as configuration settings or the measured SI channel estimates. After the operation is performed, the ARM core forwards the new RFIC canceller configuration and timing metadata back to the PC.

The compute platform’s placed-and-routed implementation characteristics on a Zynq UltraScale+ MPSoC ZCU104 Evaluation Board [61] are given in Table 4. The FPGA configures the RFIC canceller in approximately 495 μs , during

Table 4: Hardware Implementation Characteristics

Resource	Available	Utilization	Utilization
LUT (logic)	230,400	38,429	16.68%
FF	460,800	166,263	36.08%
BRAM	312	135.50	43.43%
DSP Slices	1728	174	10.07%

which the canceller is disabled and provides no analog cancellation. As this downtime is fully deterministic, transmissions can be tightly scheduled to avoid overlapping with reconfiguration; in practice, we found it sufficient to zero-pad the frames to achieve an approximately 10.8% timeslot occupancy to statistically avoid such collisions.

A typical GLC-OMP computation, selecting six taps with twenty gradient descent steps to initialize gains, takes approximately 94 ms, whereas a typical PGD computation with five gradient descent steps takes approximately 735 μs ; in both cases, the computation time scales linearly with the number of gradient descent steps. We find that we can achieve convergence within roughly ten iterations of PGD after the initial GLC-OMP stage.

Linear Digital SIC. Each packet undergoes a final digital cancellation stage to remove any residual SI, thereby allowing for the direct reconstruction of the intended receive signal. We employ a truncated Volterra series model, represented using the Toeplitz matrix-based convolution $\hat{T}\hat{h}$, where T is the Toeplitz matrix constructed from the known transmitted signal and $\hat{h}$ is the least-squares estimate of the SI channel, computed using the known pilot symbols [36].

6.2 Node-Level Cancellation

To evaluate the overall SIC performance in a practical implementation, we set up a wideband adaptive FD radio to run its full RF SIC optimization cycle while transmitting a stream of data packets. The FD radio was placed in a fixed, stationary position on a laboratory bench-top, and the room was sealed to maintain as stable an environment as possible. As an SDR-based setup allows for greater system-level insight into the performance of the canceller RFIC with data-carrying modulated waveforms, rather than purely circuit-level performance with the VNA’s CW signals, we replicated the experiments from Section 5.1 across all four bandwidths of interest (10 MHz, 25 MHz, 50 MHz, 100 MHz) centered at 1050 MHz, with the FD radio calibrated to transmit at -10 dBm.

The tap characterization procedure was performed once per bandwidth prior to experimentation. During each experiment, the SI channel was measured at each stage of the cancellation process, including the initial isolation due to the antenna interface, during each iteration of the canceller optimization, and after digital cancellation. The experiment was repeated ten times per bandwidth.




Figure 8: Residual SI channel at each iteration of the optimization over four bandwidths. In each experiment, the combination of optimal RF SIC (blue, solid) and digital SIC (red) reduced the SI to the SDR’s noise floor.




Figure 9: Average cancellation obtained at each of the four stages, compared to the noise floor target.

Figure 8 shows examples of the residual SI at each stage of the optimization for all four bandwidths; the noise floor target, determined by the average received power in the absence of a transmitted signal, is overlaid. In each of the four cases, the trajectory of the RF cancellation throughout the optimization is similar to that observed in Section 5.2. With the addition of the digital SIC, the overall cancellation becomes sufficient to reach the noise floor.³ Figure 9 shows the average results across ten trials at each bandwidth, with incremental effects at each cancellation stage; the standard deviation in RF cancellation was negligible and is omitted.

6.3 Cancellation in Mobile Environments


Whereas the previous experiments confirmed that the FD radio is capable of achieving noise-floor-level cancellation

³As can be seen in the 100 MHz experiments, two tones interfered with the experiment, appearing in the post-digital SIC residues.


in a stationary setting, it is essential to assess the system’s robustness under the dynamic conditions created by mobility. To do so, we placed the FD radio on a 1 m-high cart to act as a *mobile node*, as shown in Figure 10a. We selected a university building lobby, shown in Figure 10b, to serve as a case study for this experiment due to the richness in environmental features it provides. In particular, we identified a 27 m route (red) comprising two environments of roughly equal length: a narrow 14 m-long narrow corridor (blue) and a 13 m-long segment in a large atrium (green); for a full description of this environment, see Appendix C.

The mobile node was positioned at the start of the narrow corridor (Point 1) and was pulled at a consistent pace to the end of the open atrium (Point 2). We performed three categories of experiments in this setup: (i) no analog optimization as the mobile node moves along the route; (ii) one-shot optimization at either starting point of the traversal, with the configuration remaining static along the remainder of the route; and (iii) continuous optimization throughout the entire traversal. In all experiments, the FD radio was consistently transmitting packets over a 50 MHz bandwidth. RF SIC was continuously recorded throughout, and performance statistics were computed for the atrium and the corridor. The results of these experiments are presented in Table 5.

Although the “No Optimization” runs exhibit seemingly constant passive isolation, the underlying structure of the SI channel is in fact highly location-dependent, even at similar isolation levels. The RFIC canceller provides substantial improvements even from a single-location calibration, capturing the dominant SI structure; however, performance



(a) FD radio set up on a cart as a mobile node.



(b) Designated route, consisting of a low-ceilinged narrow corridor (blue) and a high-ceilinged open atrium (green).

Figure 10: Setup of the mobile experiment.

degrades along the path as the static configuration increasingly mismatches the evolving channel, due both to the SI's inherent dynamicity and the changing environment. This sensitivity to the initial calibration point is highlighted by the relative weakness of the “Forward One-Shot” results compared to the “Reverse One-Shot” result: the presence of a uniquely strong reflector at the optimization point (a metal door at the end of the corridor) presents a relatively anomalous environment whose ideal configuration is unsuitable for the remainder of the path, with no opportunity to recover. With continuous optimization throughout the traversal, this is no longer an issue; unsurprisingly, we see the best analog SIC performance under this regime, outperforming both “One-Shot” baselines in both environments.

We find that across all experiments, the variance is consistently higher in the corridor than in the atrium, reflecting the corridor's dense multipath environment. The nearby reflectors introduce significant delay spreads [67], requiring more time-domain cancellation paths and quick adaptation to track the abrupt SI channel changes due to the movement of the FD node. The open atrium, in contrast, presents fewer proximate reflectors and, therefore, a more spatially stable SI channel with slower evolution as the FD node progresses.

Figure 11 plots cancellation over time of the continuous optimization experiment, marking the transition point between the narrow corridor and the atrium; aggregated histograms show each region's empirical distribution. A two-sample Kolmogorov-Smirnov test [62] rejects a common underlying distribution, suggesting that analog SIC performance

Table 5: RF SIC During Mobility

Experiment	Corridor	Atrium
No Optimization	(29.9 ± 0.8) dB	(30.2 ± 0.3) dB
One-Shot (Forward)	(41.7 ± 1.5) dB	(42.4 ± 0.8) dB
One-Shot (Reverse)	(43.8 ± 2.9) dB	(45.6 ± 1.5) dB
Continuous	(45.9 ± 2.6) dB	(49.3 ± 1.4) dB

carries latent environmental information — a promising direction for environment-aware FD systems. Applying digital SIC during continuous optimization reveals that the overall system-level cancellation along the entire route was held stable at 79.9 dB with a standard deviation of 0.6 dB.

6.4 Link-Level Performance

To analyze the link-level performance of the FD radios, we experimentally evaluate average *error vector magnitude* (EVM) and *symbol error rate* (SER) as functions of link SINR. The link SINR is the average ratio between the received signal power level in the data field in the presence of the intended signal and in its absence. In HD mode, this is equivalent to the ratio of the received signal power to the Rx's noise floor; in FD mode, the residual SI degrades this metric.

EVM measures the average magnitude of the error vector between the normalized transmitted and received symbols, capturing all potential sources of signal degradation including nonlinearities, distortions, and, critically, SI. As such, it presents insight into system-level performance inaccessible when evaluating with a VNA. SER, the fraction of symbols incorrectly received, extends this concept by quantifying a radio's capacity to accurately recover data.

To demonstrate an FD link's capacity to approximately double the data rate of an equivalent HD link, we must show that an FD receiver can achieve equivalent EVM and SER performance when operating in FD and HD modes. We set up a wideband 50 MHz link between two radios fixed 2 m apart in a static environment. One node is designated as the *transmitter* and the other as the *FD receiver* that operates alternatively in FD or HD mode. In FD mode, the radios are synchronized such that the data fields overlap; to enable HD operation, the receiver's data transmission is disabled, leaving the field interference-free for the transmitter's data.

The transmitter sweeps its power while sending OFDM packets to the receiver, either in HD or FD mode; in the latter, the receiver also transmits, and converges to noise-floor-level SIC before reception begins. Data streams are randomly generated using one of three uncoded modulation schemes: binary phase-shift keying (BPSK), quadrature phase-shift keying (QPSK), and eight-phase-shift keying (8PSK).

Figure 12 shows the relationship between the receiver's EVM/SER and the link SINR across modulation schemes. EVM performance deviates by less than 0.8 dB between HD and FD modes, indicating minimal added degradation and




Figure 11: RF and overall SIC as the mobile node moves along the route, with histograms and density estimations of the RF cancellation power in each region and of the overall system cancellation power across the full route.




Figure 12: EVM and SER over SINR of the FD radio operating in HD and FD modes with different modulations.

hinting at the ability to decouple incoming packets from SI. This is confirmed by the empirical SER results, which exhibit a mean absolute log-domain error of 0.15 decades for BPSK between HD and FD modes, and 0.17 decades for QPSK and 8PSK. Though the setup is stationary, these results are applicable during mobility, as maintaining high system-level cancellation while mobile is equivalent to maintaining the EVM/SER performance in FD mode. Specifically, noise-floor-level SIC ensures that the residual received signal contains only the desired incoming signal and its impairments, equivalent to operating in HD mode; the EVM and SER performance therefore follows directly from the achieved cancellation. As shown in Figure 11, this level of cancellation is sustained throughout the traversal, confirming that the demonstrated link-level performance extends to mobile operation.

These results show that the FD radio is capable of successfully receiving and decoding packets while simultaneously transmitting over an FD link. Accounting for pilot overhead, FD delivers the same 15-slot payload per direction in one frame that HD needs two frames for — a raw $2\times$ gain. Correcting for the SER gap and for periodic RFIC reconfiguration (empirically, once every 100 ms, or a 0.5% duty cycle at 50 MHz), this reduces to approximately $1.98\times$; against an idealized HD baseline without FD pilots, the conservative gain is $1.56\times$, which improves with longer frames as overhead shrinks. At the network level, scheduling reconfiguration around idle periods can further absorb this downtime.

7 Analysis and Discussion

In this section, we analyze and discuss the results presented in Section 6, including limitations and future directions.

Cancellation Trade-Offs and Digital SIC. Compared to the results presented in Section 5.2, the SIC gain from initial isolation to optimal RF cancellation shown in Section 6.2 is lower in the practical FD radio than in the VNA-based validation using the data-driven approach, though still higher than the theory-based approach. This performance penalty reflects the trade-offs inherent in adapting the optimization algorithm for real-time, high-bandwidth hardware; we reduce the complexity and computation time, at the cost of reduced channel resolution (by reducing the number of sampling points to 242) and reduced numerical precision (by switching from 64-bit floating-point to 16-bit fixed-point).

We introduce digital SIC into the practical system to supplement the decreased RF SIC, but note that this compensation is only possible because the analog cancellation stage was sufficient to keep the residual SI power within the Rx’s dynamic range, avoiding front-end saturation and the consequent nonlinear compression that digital SIC would be unable to correct for. This was the case in both the stationary and mobile experiments: the results in Section 6.3 show that despite swings in residual SI power during continuous mobile optimization, overall cancellation was effectively flat near the noise floor; had the Rx experienced momentary saturation, digital SIC would have been unable to recover the affected packets, resulting in irrecoverable information loss.

Tx Power-Handling. A physical limitation on the maximum Tx power of the system contributed to the analog cancellation's ability to remain within the Rx's linear range throughout these experiments. The FD radio was calibrated to transmit at -10 dBm in order to account for the instantaneous power-handling capacity of the RFIC canceller (up to 7 dBm [48]) and the measured peak-to-average-power ratio (PAPR) of the OFDM-based packets (measured up to 16 dB) – a challenge absent in VNA-based evaluations, which cannot support real communication. As a result, the overall cancellation was likewise limited by this hard upper bound and the system's noise floor. This constraint also shaped our choice of modulation schemes for the link-level experiments, which were selected to maximize average Tx power while remaining within the power-handling bounds. We anticipate that higher-order MCS waveforms may see additional degradations in FD mode [36], potentially due to higher PAPR effects, which further underscores the need for advanced multi-stage cancellation mechanisms to alleviate such impacts and maximize performance.

Generalizability. Although the FD radio presented in this work operated at 1050 MHz, the underlying algorithmic contributions are not restricted to this band. Nearby frequencies are already used in a range of deployed systems, including 5G band n8 (880–960 MHz), ADS-B aeronautical navigation (978 and 1090 MHz), and ISM applications (902–928 MHz). Operation at these or other bands, including standard Wi-Fi and 5G mid-band, would require new RFIC designs; however, controlling properly for switching frequency and capacitance values, and accounting for the increased component parasitics and spurious emissions, similar switched-capacitor topologies are feasible [59], allowing the same GLC-OMP-based configuration algorithms to apply directly. Likewise, though the specific analog and digital cancellation architectures utilized here are not expected to scale efficiently to MIMO systems [7, 11], GLC-OMP itself may be adaptable to such settings, particularly for topologies composed of parallel cancellation paths with physics-constrained configuration spaces. Together, these directions suggest that the system-level contributions of this work extend well beyond the specific FD radio platform evaluated herein.

Future Directions. We plan to address the Tx power limitations by characterizing the system's nonlinearities and exploring software- and FPGA-based algorithms, including waveform shaping and digital predistortion, to improve power-handling. This will enable more complex experiments into practical FD, including over greater distances and in non-line-of-sight scenarios, and test the limits of maintaining the post-analog SIC residual within the SDR's operating range.

In addition to node-level improvements, we aim to expand into mobile link- and network-level architectures composed of multiple adaptive FD radios. This will allow us to develop

and experimentally evaluate the implications of FD operating modes for higher-layer functions, including resource scheduling and allocation, as well as the mobility-robust synchronization protocols required to enable direct mobile link-level evaluation. This will be achieved through the integration of these radios into open-access testbeds, which will supply the necessary resources (including CPUs, FPGAs, and GPUs) to deploy the nodes and provide gateway access for remote FD research to the broader community. Algorithm development, whether deployed and tested in software or on FPGA hardware, is sufficiently challenging in its own right, particularly given the real-time computation speed required at high bandwidths. We therefore aim to remove the hurdles of developing the FD systems themselves, streamlining validation through open hardware, resources, and toolkits.

8 Conclusion

In this paper, we presented our system-level design of the first wideband adaptive FD radios capable of achieving and maintaining noise-floor-level SIC in stationary and mobile environments. To achieve this, we presented a two-stage sparsity-based optimization algorithm to configure an RFIC canceller with variable gain and delay parameters, maximizing RF SIC across narrow and wide bandwidths. We evaluated this data-driven algorithm against a circuit theory-based approach and baseline approaches from the literature; confirming its improved performance, we implemented the algorithm on an FPGA as part of a practical, SDR-based FD node. Finally, we experimentally evaluated the FD radio's overall performance in stationary and mobile environments, and demonstrated that the node exhibits equivalent performance when receiving packets in either HD or FD mode.

Two key insights emerge from the work presented herein: (i) attention to device-specific impairments in RFIC canceller configuration algorithm design allows for nontrivial SIC performance gains, and (ii) intelligent system-level design for environment-aware configuration can address the challenges that dynamic scenarios, such as mobility, pose to FD advancement. Together, these results outline a path for FD radios to expand from academic works to manufacturable, field-deployable RFIC-based systems grounded in existing RF industry practices.

Acknowledgements

This work was supported in part by the DARPA WARP program, the DoD through the NDSEG Fellowship, the NSF under grants EEC-2133516, AST-2232455, CNS-2148128, and by the Federal Agency and Industry Partners as specified in the Resilient and Intelligent NextG Systems (RINGS) Program. We also thank Aditya Jolly, Ahuva Bechhofer, Eliot Flores Portillo, and Sasha Isler for their contributions.

References

- [1] 3rd Generation Partnership Project (3GPP). 2017. *LTE; Evolved Universal Terrestrial Radio Access (E-UTRA); Physical Channels and Modulation*. 3GPP TS 36.211 version 14.2.0 Release 14. European Telecommunications Standards Institute (ETSI).
- [2] In-Sung Ahn, Ji-Seong Kim, and Ho-Jin Song. 2019. Adaptive Analog Self-Interference Cancellation for In-Band Full-Duplex Wireless Communication. In *Proc. IEEE APMC'19*.
- [3] Angelos Antonopoulos, Georgios Volanis, Yichuan Lu, and Yiorgos Makris. 2019. Post-Production Calibration of Analog/RF ICs: Recent Developments and a Fully Integrated Solution. In *Proc. IEEE SMACD'19*.
- [4] Ehsan Aryafar and Alireza Keshavarz-Haddad. 2018. PAFD: Phased Array Full-Duplex. In *Proc. IEEE INFOCOM'18*.
- [5] Ehsan Aryafar, Mohammad Amir Khojastepour, Karthikeyan Sundaresan, Sampath Rangarajan, and Mung Chiang. 2012. MIDU: Enabling MIMO Full Duplex. In *Proc. ACM MobiCom'12*.
- [6] Mahmood Baraani Dastjerdi, Sanket Jain, Negar Reiskarimian, Arun Natarajan, and Harish Krishnaswamy. 2019. Analysis and Design of a Full-Duplex Two-Element MIMO Circulator-Receiver With High TX Power Handling Exploiting MIMO RF and Shared-Delay Baseband Self-Interference Cancellation. *IEEE J. Solid-State Circuits* 54, 12 (2019), 3525–3540.
- [7] Dinesh Bharadia and Sachin Katti. 2014. Full Duplex MIMO Radios. In *Proc. USENIX NSDI'14*.
- [8] Dinesh Bharadia, Emily McMillin, and Sachin Katti. 2013. Full Duplex Radios. In *Proc. ACM SIGCOMM'13*.
- [9] Bastian Bloessl, Michele Segata, Christoph Sommer, and Falko Dressler. 2013. An IEEE 802.11a/g/p OFDM Receiver for GNU Radio. In *Proc. ACM SRIF'13*.
- [10] Scott Shaobing Chen, David L. Donoho, and Michael A. Saunders. 2001. Atomic Decomposition by Basis Pursuit. *SIAM Rev.* 43, 1 (2001), 129–159.
- [11] Tingjun Chen, Mahmood Baraani Dastjerdi, Harish Krishnaswamy, and Gil Zussman. 2019. Wideband Full-Duplex Phased Array With Joint Transmit and Receive Beamforming: Optimization and Rate Gains. In *Proc. ACM MobiHoc'19*.
- [12] Tingjun Chen, Mahmood Baraani Dastjerdi, Jin Zhou, Harish Krishnaswamy, and Gil Zussman. 2019. Wideband Full-Duplex Wireless via Frequency-Domain Equalization: Design and Experimentation. In *Proc. ACM MobiCom'19*.
- [13] Tingjun Chen, Sasank Garikapati, Aravind Nagulu, Aditya Gaonkar, Manav Kohli, Igor Kadota, Harish Krishnaswamy, and Gil Zussman. 2021. A Survey and Quantitative Evaluation of Integrated Circuit-Based Antenna Interfaces and Self-Interference Cancellers for Full-Duplex. *IEEE Open J. Commun. Soc.* 2 (2021), 1753–1776.
- [14] Tingjun Chen, Jin Zhou, Nicole Grimwood, Rel Fogel, Jelena Marašević, Harish Krishnaswamy, and Gil Zussman. 2016. Demo: Full-Duplex Wireless Based on a Small-Form-Factor Analog Self-Interference Canceller. In *Proc. ACM MobiHoc'16*.
- [15] Jung Il Choi, Mayank Jain, Kannan Srinivasan, Phil Levis, and Sachin Katti. 2010. Achieving Single Channel, Full Duplex Wireless Communication. In *Proc. ACM MobiCom'10*.
- [16] Kun-Da Chu, Mohamad Katanbaf, Tong Zhang, Chenxin Su, and Jacques C Rudell. 2018. A Broadband and Deep-TX Self-Interference Cancellation Technique for Full-Duplex and Frequency-Domain Duplex Transceiver Applications. In *Proc. IEEE ISSCC'18*.
- [17] Min Keun Chung, Min Soo Sim, Jaewon Kim, Dong Ku Kim, and Chan-Byoung Chae. 2015. Prototyping Real-Time Full Duplex Radios. *IEEE Commun. Mag.* 53, 9 (2015), 56–63.
- [18] Columbia FlexCoN Project. 2026. Code and Documentation. <https://github.com/Wimnet/flexicon-wideband-adaptive>.
- [19] Jogesh Chandra Dash and Debdeep Sarkar. 2022. A Colinearly Polarized Full-Duplex Antenna With Extremely High Tx–Rx Isolation. *IEEE Antennas Wireless Propag. Lett.* 21, 12 (2022), 2387–2391.
- [20] Yongqian Du, Jiahao Chen, Guifang Li, Kai Han, Tianxiang Qu, and Shibin Liu. 2022. Integrated Self-Interference Suppression for Single Antenna Full-Duplex Without Circulator. In *Proc. IEEE NEWCAS'20*.
- [21] Michael Elad. 2010. *Sparse and Redundant Representations: From Theory to Applications in Signal and Image Processing*. Springer.
- [22] Ali Ershadi and Kamran Entesari. 2020. A 0.5-to-3.5-GHz Full-Duplex Mixer-First Receiver With Cartesian Synthesized Self-Interference Suppression Interface in 65-nm CMOS. *IEEE Trans. Microw. Theory Techn.* 68, 6 (2020), 1995–2010.
- [23] Ettus Research. 2022. USRP Hardware Driver (UHD) Software 4.2. <https://github.com/EttusResearch/uhd>.
- [24] Evan Everett, Clayton Shepard, Lin Zhong, and Ashutosh Sabharwal. 2016. SoftNull: Many-Antenna Full-Duplex Wireless via Digital Beamforming. *IEEE Trans. Wireless Commun.* 15, 12 (2016), 8077–8092.
- [25] Robin Garg, Sanket Jain, and Arun S. Natarajan. 2022. Wideband, Millimeter Wave Domain SI Canceling (>50dB) In-Band Full-Duplex Circulator Receiver. *IEEE Access* 10 (2022), 37953–37966.
- [26] Sasank Garikapati, Aravind Nagulu, Igor Kadota, Mostafa Essawy, Tingjun Chen, Shibo Wang, Tanvi Pande, Arun S. Natarajan, Gil Zussman, and Harish Krishnaswamy. 2024. Full-Duplex Receiver With Wideband, High-Power RF Self-Interference Cancellation Based on Capacitor Stacking in Switched-Capacitor Delay Lines. *IEEE J. Solid-State Circuits* 59, 7 (2024), 2105–2120.
- [27] Sastry Garimella, Sasank Garikapati, Aravind Nagulu, Igor Kadota, Alfred Davidson, Gil Zussman, and Harish Krishnaswamy. 2023. Frequency-Domain-Equalization-Based Full-Duplex Receiver With Passive-Frequency-Shifting N-Path Filters Achieving >53 dB SI Suppression Across 160 MHz BW. In *Proc. IEEE RFIC'23*.
- [28] GNU Radio. 2019. GNU Radio 3.8. <http://gnuradio.org/>.
- [29] Malik Muhammad Usman Gul, Xiaoli Ma, and Sungeun Lee. 2015. Timing and Frequency Synchronization for OFDM Downlink Transmissions Using Zadoff-Chu Sequences. *IEEE Trans. Wireless Commun.* 14, 3 (2015), 1716–1729.
- [30] Hiroshi Harada, Keiichi Mizutani, Takeshi Matsumura, Toyoyuki Kato, and Ken Shioiri. 2022. Development of Full-Duplex Cellular System for Beyond 5G and 6G Systems. In *Proc. IEEE PIMRC'22*.
- [31] Mayank Jain, Jung Il Choi, Taemin Kim, Dinesh Bharadia, Siddharth Seth, Kannan Srinivasan, Philip Levis, Sachin Katti, and Prasun Sinha. 2011. Practical, Real-Time, Full Duplex Wireless. In *Proc. ACM MobiCom'11*.
- [32] Edward Kassem, Roman Marsalek, and Jiri Blumenstein. 2018. Frequency Domain Zadoff-Chu Sounding Technique for USRPs. In *Proc. IEEE ICT'18*.
- [33] Seiran Khaledian, Frahad Farzami, Besma Smida, and Danilo Erricolo. 2018. Inherent Self-Interference Cancellation at 900 MHz for In-Band Full-Duplex Applications. In *Proc. IEEE WAMICON'18*.
- [34] Eric A.M. Klumperink, Hugo J. Westerveld, and Bram Nauta. 2017. N-Path Filters and Mixer-First Receivers: A Review. In *Proc. IEEE CICC'17*.
- [35] Manav Kohli, Mahmood Baraani Dastjerdi, Jin Zhou, Ivan Seskar, Harish Krishnaswamy, Gil Zussman, and Tingjun Chen. 2025. Design and Testbed Deployment of Frequency-Domain Equalization Full Duplex Radios. *IEEE Trans. Wireless Commun.* 24, 5 (2025), 4377–4391.
- [36] Manav Kohli, Tingjun Chen, Mahmood Baraani Dastjerdi, Jackson Welles, Ivan Seskar, Harish Krishnaswamy, and Gil Zussman. 2021. Open-Access Full-Duplex Wireless in the ORBIT and COSMOS Testbeds. *Computer Networks* 199 (2021), 108420.
- [37] Kenneth E. Kolodziej, Aidan U. Cookson, and Bradley T. Perry. 2020. In-Band Full-Duplex RF Canceller Tuning Using Adaptive Learning

- Rate Functions. In *Proc. IEEE RWS'20*.
- [38] Kenneth E. Kolodziej, Joseph G. McMichael, and Bradley T. Perry. 2016. Multitap RF Canceller for In-Band Full-Duplex Wireless Communications. *IEEE Trans. Wireless Commun.* 15, 6 (2016), 4321–4334.
- [39] Kenneth E. Kolodziej, Bradley T. Perry, and Jeffrey S. Herd. 2019. In-Band Full-Duplex Technology: Techniques and Systems Survey. *IEEE Trans. Microw. Theory Techn.* 67, 7 (2019), 3025–3041.
- [40] Dani Korpi, Joose Tamminen, Matias Turunen, Timo Huusari, Yang-Seok Choi, Lauri Anttila, Shilpa Talwar, and Mikko Valkama. 2016. Full-Duplex Mobile Device: Pushing the Limits. *IEEE Commun. Mag.* 54, 9 (2016), 80–87.
- [41] Harish Krishnaswamy and Gil Zussman. 2016. 1 Chip, 2x the Bandwidth. *IEEE Spectrum* 53, 7 (2016), 38–54.
- [42] Leo Laughlin, Chunqing Zhang, Mark A. Beach, Kevin A. Morris, and John L. Haine. 2016. Passive and Active Electrical Balance Duplexers. *IEEE Trans. Circuits Syst. II, Exp. Briefs* 63, 1 (2016), 94–98.
- [43] Alon S. Levin, Eliot Samuel Flores Portillo, Sasank Garikapati, Ahuva Bechhofer, Bo Zhang, Manav Kohli, Igor Kadota, Harish Krishnaswamy, Mingoo Seok, and Gil Zussman. 2024. Demo: Achieving Self-Interference Cancellation Across Different Environments. In *Proc. ACM MobiCom'24*.
- [44] Alon S. Levin, Igor Kadota, Sasank Garikapati, Bo Zhang, Aditya Jolly, Manav Kohli, Mingoo Seok, Harish Krishnaswamy, and Gil Zussman. 2023. Demo: Experimentation With Wideband Real-Time Adaptive Full-Duplex Radios. In *Proc. ACM SIGCOMM'23*.
- [45] Alon S. Levin, Manav Kohli, Igor Kadota, Tingjun Chen, Sasank Garikapati, Aravind Nagulu, Mahmood Baraani Dastjerdi, Jin Zhou, Ivan Seskar, Harish Krishnaswamy, and Gil Zussman. 2024. Enabling Integrated Circuit-Based Full-Duplex Wireless in GNU Radio. In *Proc. GRCon'24*.
- [46] Mohyee Mikhemar, Hooman Darabi, and Asad A. Abidi. 2013. A Multiband RF Antenna Duplexer on CMOS: Design and Performance. *IEEE J. Solid-State Circuits* 48, 9 (2013), 2067–2077.
- [47] Carl W. Morgenstern, Yu Rong, Andrew Herschfeld, Alyosha C. Molnar, Alyssa B. Apsel, David G. Landon, and Daniel W. Bliss. 2022. Tuning Algorithm to Compensate for Hardware Errors of IBFD Analog Device. In *Proc. Asilomar'22*.
- [48] Aravind Nagulu, Sasank Garikapati, Mostafa Essawy, Igor Kadota, Tingjun Chen, Arun Natarajan, Gil Zussman, and Harish Krishnaswamy. 2021. Full-Duplex Receiver With Wideband Multi-Domain FIR Cancellation Based on Stacked-Capacitor, N-Path Switched-Capacitor Delay Lines Achieving >54dB SIC Across 80MHz BW and >15dBm TX Power-Handling. In *Proc. IEEE ISSCC'21*.
- [49] Aravind Nagulu, Negar Reiskarimian, Tingjun Chen, Sasank Garikapati, Igor Kadota, Tolga Dinc, Sastry Lakshmi Garimella, Manav Kohli, Alon S. Levin, Gil Zussman, and Harish Krishnaswamy. 2024. Doubling Down on Wireless Capacity: A Review of Integrated Circuits, Systems, and Networks for Full Duplex. *Proc. IEEE* 112, 5 (2024), 405–432.
- [50] John C. L. Ng, Khaled Ben Letaief, and Ross D. Murch. 1998. Complex Optimal Sequences With Constant Magnitude for Fast Channel Estimation Initialization. *IEEE Trans. Commun.* 46, 3 (1998), 305–308.
- [51] Prathap Valale Prasannakumar, Mohamed A. Elmansouri, and Dejan S. Filipovic. 2017. Wideband Decoupling Techniques for Dual-Polarized Bi-Static Simultaneous Transmit and Receive Antenna Subsystem. *IEEE Trans. Antennas Propag.* 65, 10 (2017), 4991–5001.
- [52] Bozidar Radunovic, Dinan Gunawardena, Peter Key, Alexandre Proutiere, Nikhil Singh, Vlad Balan, and Gerald Dejean. 2010. Re-thinking Indoor Wireless Mesh Design: Low Power, Low Frequency, Full-Duplex. In *Proc. IEEE WIMESH'10*.
- [53] Negar Reiskarimian, Jin Zhou, and Harish Krishnaswamy. 2017. A CMOS Passive LPTV Nonmagnetic Circulator and Its Application in a Full-Duplex Receiver. *IEEE J. Solid-State Circuits* 52, 5 (2017), 1358–1372.
- [54] Ashutosh Sabharwal, Philip Schniter, Dongning Guo, Daniel W. Bliss, Sampath Rangarajan, and Risto Wichman. 2014. In-Band Full-Duplex Wireless: Challenges and Opportunities. *IEEE J. Sel. Areas Commun.* 32, 9 (2014), 1637–1652.
- [55] Vaibhav Singh, Susnata Mondal, Akshay Gadre, Milind Srivastava, Jeyanandh Paramesh, and Swarun Kumar. 2020. Millimeter-Wave Full Duplex Radios. In *Proc. ACM MobiCom'20*.
- [56] Besma Smida, Ashutosh Sabharwal, Gabor Fodor, George C. Alexandropoulos, Himal A. Suraweera, and Chan-Byoung Chae. 2023. Full-Duplex Wireless for 6G: Progress Brings New Opportunities and Challenges. *IEEE J. Sel. Areas Commun.* 41, 9 (2023), 2729–2750.
- [57] Joose Tamminen, Matias Turunen, Dani Korpi, Timo Huusari, Yang-Seok Choi, Shilpa Talwar, and Mikko Valkama. 2016. Digitally-Controlled RF Self-Interference Canceller for Full-Duplex Radios. In *Proc. EUSIPCO'16*.
- [58] Joel A. Tropp. 2004. Greed Is Good: Algorithmic Results for Sparse Approximation. *IEEE Trans. Inf. Theory* 50, 10 (2004), 2231–2242.
- [59] Mykhailo Tymchenko, Dimitrios Sounas, Aravind Nagulu, Harish Krishnaswamy, and Andrea Alu. 2019. Quasi-Electrostatic Wave Propagation Beyond the Delay-Bandwidth Limit in Switched Networks. *Physical Review X* 9, 3 (2019), 031015.
- [60] Xiyao Wang, Wenquan Che, Wanchen Yang, Wenjie Feng, and Lizheng Gu. 2017. Self-Interference Cancellation Antenna Using Auxiliary Port Reflection for Full-Duplex Application. *IEEE Antennas Wireless Propag. Lett.* 16 (2017), 2873–2876.
- [61] Xilinx. 2018. *ZCU104 Evaluation Board User Guide*. UG1267 (v1.1). <https://docs.amd.com/v/u/en-US/ug1267-zcu104-eval-bd>.
- [62] Guowei Zhang, Xiaodong Wang, Ying-Chang Liang, and Ju Liu. 2010. Fast and Robust Spectrum Sensing via Kolmogorov-Smirnov Test. *IEEE Trans. Commun.* 58, 12 (2010), 3410–3416.
- [63] Tong Zhang, Chenxin Su, Ali Najafi, and Jacques Christophe Rudell. 2018. Wideband Dual-Injection Path Self-Interference Cancellation Architecture for Full-Duplex Transceivers. *IEEE J. Solid-State Circuits* 53, 6 (2018), 1563–1576.
- [64] Zhongshan Zhang, Keping Long, Athanasios V. Vasilakos, and Lajos Hanzo. 2016. Full-Duplex Wireless Communications: Challenges, Solutions, and Future Research Directions. *Proc. IEEE* 104, 7 (2016), 1369–1409.
- [65] Jin Zhou, Anandaroop Chakrabarti, Peter Kinget, and Harish Krishnaswamy. 2014. Low-Noise Active Cancellation of Transmitter Leakage and Transmitter Noise in Broadband Wireless Receivers for FDD/Co-Existence. *IEEE J. Solid-State Circuits* 49, 12 (2014), 1–17.
- [66] Jin Zhou, Tsung-Hao Chuang, Tolga Dinc, and Harish Krishnaswamy. 2015. Integrated Wideband Self-Interference Cancellation in the RF Domain for FDD and Full-Duplex Wireless. *IEEE J. Solid-State Circuits* 50, 12 (2015), 3015–3031.
- [67] Jin Zhou and Harish Krishnaswamy. 2018. System-Level Analysis of Phase Noise in Full-Duplex Wireless Transceivers. *IEEE Trans. Circuits Syst. II, Exp. Briefs* 65, 9 (2018), 1189–1193.

A Background: The RFIC Canceller Circuit

In this appendix, we expand on the operating principles of the RFIC canceller summarized in Section 3. The RFIC canceller is integrated into the frontend of the FD radio on a signal path parallel to that of the antenna interface, as shown in Figure 13, and acts as a filter aiming to match the SI channel experienced at the antenna interface. The signal received at the antenna, consisting of the desired signal and the SI signal,

is combined with the Tx signal filtered through the RFIC; in an ideal outcome, the RFIC canceller perfectly inverts the SI channel, resulting in complete destructive interference of the SI signal and fully isolating the desired signal.

Circuit-Level Overview. The RFIC canceller is integrated into an FD receiver IC implemented in a standard 64 nm CMOS process. The receiver consists of a DC–1 GHz partial-noise-canceling low-noise transimpedance amplifier (LNTA), followed by a 4-phase I/Q direct down-conversion mixer and a baseband transimpedance amplifier. Tx power is capacitively coupled into the RFIC using digitally programmable off-chip capacitors. The RFIC consists of source-follower-based buffers at the input, which drive 16 individual delay taps, realized using switched-capacitor-based delays [59].

The individual taps fall into three tap classes as shown in Table 2: (i) one zero-delay tap in which the input and output are directly connected, (ii) five low-delay taps realized using eight parallel time-interleaved switched-capacitor-based delay paths providing up to 2 ns delay with a resolution of 250 ps across a 1 GHz bandwidth when clocked at a 500 MHz frequency, and (iii) ten high-delay taps realized using 32 parallel time-interleaved switched-capacitor-based delay paths providing up to 8 ns delay with a resolution of 250 ps across a 1 GHz bandwidth when clocked at a 500 MHz frequency. The high-delay taps benefit from a lower clock frequency to obtain higher delay and utilize more interleaved paths to preserve the bandwidth of operation. However, owing to the parasitic capacitances due to the large number of interleaved paths, these delay taps have increased loss compared to the low-delay taps, especially at higher frequencies. The RFIC canceller utilizes a capacitor-stacking-based delay element, which provides passive voltage gain and compensates for the parasitic losses of the delay taps.

The individual tap outputs are connected to the receiver's novel LNTA canceller, which performs FIR weighting, summation of the taps, and signal injection into the Rx to achieve cancellation. Thus, the RFIC canceller realizes an adaptive FIR filter using parallel capacitor-stacking-based delay elements with tunable gains. These delay elements directly interface with the receiver's LNTA to inject the desired current required for cancellation at the Rx. Further, the receiver's LNTA can also be reconfigured to behave as an output buffer for the RFIC canceller, enabling the chip to operate as a standalone canceller for use with off-chip receivers.

Capacitor-Stacking-Based Delay Element. The delay element uses a time-interleaved multi-path switched-capacitor network, as shown in Figure 14. In this, the input signal is periodically sampled onto a capacitor with a clocking frequency f_s and a pulse width of T_s/N , where N is the number of interleaved paths. The capacitor's stored sample is sensed at the output through a switch controlled by a delayed version of the sampling clock. This delay between the sampling




Figure 13: Block diagram of the RFIC canceller.




Figure 14: Mechanism of the delay-generating circuits.

and the sensing clocks provides a true-time delay from the input to the output, and can be easily programmed by varying the clock staggering. Multiple such switched capacitors are interleaved in time such that at any instant, the input signal is being sampled onto one of the capacitors, and one of the capacitors is being sensed at the output. The maximum possible amount of delay is the amount of time between two successive charging cycles of the same capacitor, which is approximately T_s . The bandwidth is inversely proportional to the pulse width and hence is directly proportional to $N \times f_s$. Thus, decreasing the clocking frequency increases T_s and thereby the maximum delay, requiring N to be proportionally increased in order to maintain the same bandwidth.

Switch parasitics and dead time between the multi-phase clocks result in increased loss, which compromises the power handling and noise performance of the delay elements. This work utilizes a capacitor-stacking approach to achieve passive voltage gain in the delay element, preserving the linearity and noise performance. In this scheme, during the sampling phase, the input signal is sampled onto multiple parallel capacitors instead of a single capacitor. During the sensing phase, the individual capacitors are stacked over each other in series and are connected to the output node such that the output voltage equals the sum of the voltages

on the individual capacitors. This provides a passive voltage gain equal to the number of stacked capacitors.

B Derivations and Algorithms

In this appendix, we expand on Section 4 by providing more details on the algorithms and mathematical derivations that support the problem formulations and algorithmic solutions; furthermore, we expand on the implementation of the baseline algorithms employed in Section 5.

B.1 Proposed Algorithms

Theory-Based Formulation & Solution. Assuming that the RFIC canceller's true performance closely matches the ideal analytical model in Equation (2), the RFIC canceller's frequency response can be modelled as the sum of the individual tap responses, namely,

$$H_{\text{RFIC}}(f, \theta) = \sum_{n=0}^{15} H_n(f, \theta_n). \quad (15)$$

In turn, the optimization problem in Equation 1 becomes:

$$\theta^* := \arg \min_{\theta \in \Theta} \left\| H_{\text{SI}}(f) - \sum_{n=0}^{15} \alpha_n e^{-j2\pi f \tau_n} \right\|^2. \quad (16)$$

Projected Gradient Descent (PGD), shown in Algorithm 2, can be used to solve this optimization, using the iterative update function

$$\theta_{t+1} \leftarrow \text{Proj}[\theta_t - \eta_t \nabla g(\theta_t)], \quad (17)$$

where η_t is the learning rate, $g(\theta)$ is the objective function in Equation 1, and $\text{Proj}[\theta]$ is an operator that projects an out-of-bounds configuration θ onto the nearest in-bound limit point (e.g., $\alpha_n \in [0, 1]$). The gradient of the objective

Algorithm 2 Projected Gradient Descent (PGD)

Input: SI channel $H_{\text{SI}}(f)$, configuration θ

```

1: for  $t = 1$  to  $t_{\text{max}}$  do
2:    $\eta_t \leftarrow 1$ 
3:   Compute  $\text{Res}(f, \theta)$   $\triangleright$  Equation 1
4:   Compute  $\nabla g(\theta)$   $\triangleright$  Equation 18
5:   repeat  $\triangleright$  Backtracking line search
6:      $G_t(\theta) \leftarrow (\theta - \text{Proj}[\theta - \eta_t \nabla g(\theta)]) / \eta_t$ 
7:     LHS  $\leftarrow g(\text{Proj}[\theta - \eta_t \nabla g(\theta)])$ 
8:     RHS  $\leftarrow g(\theta) - \eta_t \nabla g(\theta)^\top G_t(\theta) + \frac{\eta_t}{2} \|G_t(\theta)\|_2^2$ 
9:      $\eta_t \leftarrow \eta_t / 2$ 
10:  until LHS < RHS
11:   $\theta \leftarrow \text{Proj}[\theta - \eta_t \nabla g(\theta)]$ 
12:  if  $\|\nabla g(\theta)\|^2 < \epsilon$  then
13:    break
14:  end if
15: end for
16: return Quantize( $\theta$ )
```

function is determined analytically as a 32-element vector:

$$\begin{aligned} \nabla g(\theta) &= \begin{bmatrix} \frac{\partial g(\theta)}{\partial \alpha_n} \Big|_{n=0}^{15} \\ \frac{\partial g(\theta)}{\partial \tau_n} \Big|_{n=0}^{15} \end{bmatrix} \\ &= \begin{bmatrix} -2 \text{Re} [\exp\{j2\pi f \tau_n\} \text{Res}(f, \theta_n)] \Big|_{n=0}^{15} \\ 4\pi \alpha_n f \text{Im} [\exp\{j2\pi f \tau_n\} \text{Res}(f, \theta_n)] \Big|_{n=0}^{15} \end{bmatrix}. \end{aligned} \quad (18)$$

Backtracking line search is used to determine the learning rate for each update step, and the algorithm is performed for a fixed number of iterations or until a stopping criterion on the gradient size is achieved.

Iterative Gain Adjustment. Algorithm 3 is a modification of Algorithm 2, and serves as the iterative adjustment phase of the data-driven approach from Section 4.3. In particular, this modification employs the substitutions described in Equations 13 and 14.

B.2 Baseline Algorithms

Cubic Polynomial Modeling. In [26], the authors propose a greedy algorithm that makes use of a cubic tap model to perform RFIC canceller optimization. In this method, each measured (tap, delay) tuple $\hat{H}_n(f, \tau_n)$ is fit to a cubic polynomial model, as shown in Figure 2, with a general form

$$H_n^{\text{cub}}(f) = e^{-j2\pi f \tau_n} e^{\hat{a} + \hat{b}f\tau_n + \hat{c}f^2\tau_n^2 + \hat{d}f^3\tau_n^3}. \quad (19)$$

A nonlinear least-squares solver is employed to compute the complex-valued coefficients $\hat{a}$, $\hat{b}$, $\hat{c}$, and $\hat{d}$. The resulting cubic tap models are then used in a greedy algorithm followed by iterative gradient descent-based optimization to determine the optimal SIC-maximizing canceller configuration. Two important differences with respect to our approach are:

Algorithm 3 Modified PGD

Input: SI channel $H_{\text{SI}}(f)$, gain vector $\mathbf{A}$

```

1: for  $t = 1$  to  $t_{\text{max}}$  do
2:    $\eta_t \leftarrow 1$ 
3:   Compute  $\nabla g(\mathbf{A})$   $\triangleright$  Equation 13
4:   repeat  $\triangleright$  Backtracking Line Search
5:      $G_t(\mathbf{A}) \leftarrow (\mathbf{A} - \text{Proj}[\mathbf{A} - \eta_t \nabla g(\mathbf{A})]) / \eta_t$ 
6:     LHS  $\leftarrow g(\text{Proj}[\mathbf{A} - \eta_t \nabla g(\mathbf{A})])$ 
7:     RHS  $\leftarrow g(\mathbf{A}) - \eta_t \nabla g(\mathbf{A})^\top G_t(\mathbf{A}) + \frac{\eta_t}{2} \|G_t(\mathbf{A})\|_2^2$ 
8:      $\eta_t \leftarrow \eta_t / 2$ 
9:   until LHS < RHS
10:   $\mathbf{A} \leftarrow \text{Proj}[\mathbf{A} - \eta_t \nabla g(\mathbf{A})]$ 
11:  if  $\|\nabla g(\mathbf{A})\|^2 < \epsilon$  then
12:    break
13:  end if
14: end for
15: return Quantize( $\mathbf{A}$ )
```

(i) the cubic model for the measured tap responses; and (ii) the greedy algorithm (which is not based on OMP).

Straddled-Delay Selection. The approach proposed in [8] relies on an intentional delay selection based on the measured principal delay of the SI channel. The system in [8] utilizes sixteen fixed delay lines implemented using varying-length micro-strip traces with individually programmable attenuators. The delays were selected to form an equally-spaced range straddling the principal delay of the SI leakage path through a circulator. This results in lagging and leading copies of the SI signal that can be linearly combined to reconstruct the original, thereby providing cancellation. To emulate this, we measured the average group delays of each tap-delay tuple, and selected delay settings to achieve the most uniform spacing around the observed SI delay. Keeping these delay settings fixed, we solve a linear least squares problem to determine initial gain values, and then iteratively adjust with gradient descent.

C Implementation Details

In this appendix, we provide implementation and deployment details of the wideband, adaptive FD radio system described in Section 6, including a more in-depth discussion of the SDR and host PC configuration used to sustain reliable high-bandwidth operation, a description of the FPGA’s host-communication protocol and operating modes, and the specifications of the physical environment used in the mobility experiments presented in Section 6.3. We note that many iterations of hardware and software optimizations have been performed to achieve the presented results; we discuss these optimizations below.

C.1 FD Radio Configuration

Software-Defined Radio. Each FD radio utilizes a USRP X310 SDR with UBX-160 daughterboards, capable of operating within the RFIC canceller’s operating range with up to 160 MHz of instantaneous bandwidth. We employ two COTS swivel whip antennas (Siretta Delta 6A), with operating frequency ranges spanning the ISM 915 MHz band and above, as the antenna interface to the FD radio.

The SDRs are synchronized by a single external reference clock and PPS signal, ensuring frequency lock and timing accuracy. The SDR streams 16-bit complex (IQ) samples to and from the host PC over a direct 10 GbE LAN interface. All communication and control with the SDR is facilitated by the USRP Hardware Driver (UHD) software package [23].

The SDR expects the bidirectional sample flow, to and from the PC, to be set at a constant rate. At high data rates, the host PC may not produce or consume data fast enough, dropping packets in the process and resulting in under run or

overflow errors, respectively. Occasional errors are generally non-fatal: recovery can be performed via software-based resynchronization, which accounts for the packets’ shifted positions and any resulting sub-sample timing offsets. However, multiple errors in a short span will result in a full throttling of the system, freezing all processes and requiring a full system restart. We found that even moderate-performance PCs failed to reliably transmit and receive at a 50 MHz sampling rate due to compounding errors; therefore, to minimize the likelihood of such errors, we built a custom host PC with optimized hardware capable of handling these data streams. **Host PC.** We use a 24-core Intel i9-13900K CPU, and the frequency governor was set to performance mode across all cores to prevent clock scaling during operation. We use 4×16 GB of DDR5 RAM. The OS, Ubuntu 22.04.5 LTS, is installed on a 1 TB Gen 5 NVMe SSD, and bulk data is stored on a 4 TB external HDD. All experiments are performed using GNU Radio 3.8 [28], with custom control and DSP blocks implemented in C++ [45].

Several network-level optimizations were applied to reduce the risk of packet stalling between the host PC and the SDR. The motherboard’s onboard 10 GbE NIC was configured with an MTU of 9000 bytes (jumbo frames) to reduce per-packet overhead at high data rates. Both the send and receive Ethernet ring buffers were increased to 4096 to help prevent flow control errors, and both the send and receive network socket buffers were increased to 32 MB.

C.2 FPGA Operations

We developed a custom C++ library that defines the data frames sent between the host PC and the FPGA, which are connected over UART. The host-to-FPGA frame consists of three primary fields: an *operation flag*, which defines the operational mode and hyperparameter settings for the FPGA to use; a *configuration field*, representing the gains and delays for a six-tap configuration; and a *channel estimate field*, representing a 242-point complex $H_{SI}(f)$ measurement. The FPGA-to-host response consists of a *configuration field*, indicating the results of the requested operation and the settings loaded onto the canceller, and a *timing metadata field* containing runtime information.

The ARM core of the FPGA runs a control program with three operational modes: *manual system configuration*, *GLC-OMP-based computation*, and *PGD-based computation*.

In the system configuration mode, the FPGA only makes use of the configuration field. The ARM core decodes the data and writes the configuration bits to Mem_{cfg} . `CanConfig` then reads the data from Mem_{cfg} and configures the RFIC canceller. This operation is used for initial configuration before optimization, the tap measurement process, and manual configuration.

In the GLC-OMP-based computation mode, the FPGA makes use of the operation flag and the channel estimate field. The ARM core loads the hyperparameter settings and forwards the channel estimate to the ConfigOpt block. The FPGA then executes the GLC-OMP-based optimization algorithm and selects the optimal taps and delays to enable, along with their respective gain values. As before, this configuration is sent to the Mem_{cfg} block, from which the CanConfig block fetches the data and configures the canceller.

In the PGD-based computation mode, the FPGA makes use of all fields. The ARM core decodes the configuration field, and sets the indicated (tap, delay) settings and their assigned gains as the starting point for the PGD algorithm. Using the requested hyperparameter settings, the ConfigOpt block processes the channel estimate to compute new gain values for the taps. As in the other modes, the CanConfig block fetches the new configuration from the Mem_{cfg} block and configures the canceller.

C.3 Mobile Experiment Environment

The 27 m route used for the mobility experiment, shown in red in Figure 10b, is comprised of two environments of roughly equal length. The first is a 14 m-long corridor (shown in blue) with a low ceiling (2.7 m), narrow width (1.6 m), and various metal doors, wall recesses, and conduit runs along the path. The second is a 13 m-long segment through a large atrium (shown in green) with a high ceiling (5.4 m). The first four meters are bounded by a nearby wall and a glass doorway 5 m away; the remaining 9 m-path opens into a wide region 5 m across. A floor-to-ceiling column acts as the transition between the regions, as demarcated on the map.