

On Self-interference Suppression Methods for Low-complexity Full-duplex MIMO

Alexios Balatsoukas-Stimming, Pavle Belanovic, Konstantinos Alexandris, Andreas Burg
 École Polytechnique Fédérale de Lausanne (EPFL)
 Telecommunications Circuits Laboratory (TCL)

Abstract—Full-duplex wireless communication offers improved spectral efficiency, as well as more efficient relaying and medium access, but requires suppression of self-interference. In this paper we analyze the existing methods for active RF suppression and use the “Rice architecture” for its low complexity and favorable scaling when applied to multi-antenna systems. We analyze the effects of the different sources of self-interference and quantify the potential for further suppression (genie-aided suppression). Our single-chain implementation using a circulator achieves -48 dB of active RF suppression, but only -66 dB of total suppression in the analog domain. On the other hand, our single-chain implementation using separate antennae reaches -85 dB of total analog suppression, thus reducing the self-interference to the noise floor. Extending these setups, we present a low complexity implementation of a 2×2 full-duplex MIMO node, which achieves even higher suppression than the single-chain counterparts.

I. INTRODUCTION

Full-duplex (FD) wireless communication has recently gained much attention. It promises significant improvement in spectral efficiency compared to the traditional half-duplex (HD) mode. This efficiency gain exists in both single-input single-output (SISO) as well as multiple-input multiple-output (MIMO) wireless systems. In addition, FD also offers more efficient relaying and medium access control.

However, a key weakness of FD is the self-interference from the node’s own transmitted signal. This self-interference is orders of magnitude above the noise floor and the desired signal from any distant node in the network. Hence, for a node to operate successfully in FD mode, the node must implement self-interference suppression. The suppression must be high enough to push the power of the self-interference down to the noise floor, so that it can be treated in the same way as the usual additive receiver noise.

A. Stages of suppression

Two basic types of self-interference suppression exist: passive and active. Passive suppression simply refers to the natural attenuation that the self-interference signal experiences as it propagates from the transmit to the receive chain. Active suppression refers to any additional signal processing carried out on the received signal, specifically aimed at reducing self-interference, such as the subtraction of a generated cancelation signal.

Typically, no single mechanism of active self-interference suppression is sufficiently effective to bring the power of the self-interference down to the noise floor. Hence, a sequence of active suppression mechanisms must be implemented, each at a different stage of processing the received signal. The usual stages are the electromagnetic (EM), radio frequency (RF), analog baseband (BB), and digital. Self-interference suppression in the analog stages is critical, before the signal reaches the analog to digital converter (ADC), in order to

Corresponding author: Alexios Balatsoukas-Stimming (alexios.balatsoukas@epfl.ch). This work was supported in part by the Swiss National Science Foundation project FORTE.

make maximal use of the ADC’s limited dynamic range. Although substantial suppression in the EM domain is achievable, it can also impair the far field operation of the transceiver [1]. Suppression in the BB domain may be difficult to implement in practice using off-the-shelf radios, due to the need to access the received signal after the local oscillator, but before the ADC. Hence, recent contributions mostly focused on active RF and digital suppression.

In this paper, we also consider the combination of active RF and digital self-interference suppression. At both stages we reproduce the self-interference signal, and subtract it from the received signal. It is important to note that all the components of the self-interference signal are accessible to us (i.e., cancelable) at both stages. In other words, there is no suppression that can be performed in the digital domain, that could not have been performed earlier in the RF domain. In summary, the first pass of suppression is applied in the analog domain at the RF stage, and the digital suppression that follows is a repetition of the same process, as a second clean-up pass, that only corrects the imperfections of the first pass.

B. Related work

Although several different methods of implementing active RF suppression have been shown, they all fall into just two general categories. In all cases, the goal is to recreate the self-interference signal (by transforming the known transmitted signal) and subtract it from the received signal. The Stanford architecture [1]–[3] focuses on recreating the self-interference signal by directly transforming a copy of the *analog RF* signal using analog components. On the other hand, the Rice architecture [4], [5] advocates recreating the self-interference in the *digital baseband* domain, followed by a separate RF chain to generate the RF cancelation signal. Each approach has its benefits and weaknesses.

II. LIMITATIONS OF ACTIVE RF SUPPRESSION

In order to identify the performance limitations of each approach, we first provide a simple model for both techniques.

A. Stanford architecture

The structure of the Stanford architecture is shown in Figure 1. Let $s(t)$ be the own transmitted signal, $x(t)$ be the desired remote signal, $e_{tx}(t)$ be the transmitter noise, $n(t)$ be the receiver thermal noise, and $q(t)$ be the quantization noise, where t is time. Let also $h_{tx}(\tau)$ be the impulse response of the wireless channel experienced by the self-interference signal, $\hat{h}_{tx}(\tau)$ be the reconstruction of $h_{tx}(\tau)$ produced by the RF suppression method in question, and $h_x(\tau)$ be the impulse response of the wireless channel experienced by the remote signal, where τ is also a time variable. We represent all of these signals as random, strictly stationary, continuous-time, complex, baseband signals. The received signal, r_s , is given by:

$$r_s = \underbrace{(s + e_{tx}) * (h_{tx} - \hat{h}_{tx})}_{\text{self-interference, } z_s} + x * h_x + n + q, \quad (1)$$




Fig. 1. Stanford architecture, the suppression circuitry is highlighted in gray.

where $*$ denotes the convolution operator. For simplicity we omit the time variables t and τ . Now, the self-interference signal after active RF suppression, z_s , is:

$$z_s = (s + e_{tx}) * \tilde{h}_s, \quad (2)$$

where $\tilde{h}_s = h_{tx} - \hat{h}_{tx}$ is the residual channel, or the error in estimating and replicating h_{tx} . Let the active RF suppression, β_s , be the ratio of the interference power after suppression, to the interference power before suppression. Of course, ideally we would have $\beta_s = 0$. For any given h_s and h_{tx} , we see that:

$$\beta_s = \frac{E \left[|(s + e_{tx}) * \tilde{h}_s|^2 \right]}{E \left[|(s + e_{tx}) * h_{tx}|^2 \right]} = \frac{\int_{-\infty}^{\infty} |\tilde{h}_s|^2}{\int_{-\infty}^{\infty} |h_{tx}|^2}, \quad (3)$$

where $E[\cdot]$ is the expectation operator and the integral is taken over the omitted time variable τ . Hence, suppression in this architecture is limited by how well we are able to estimate and replicate the self-interference wireless channel.

While h_{tx} is unknown and time-varying, it is possible to tune the analogue circuitry, as shown in [3], to maximize suppression, achieving up to -45dB. However, the need to replicate a multi-path h_{tx} with several energy components at different delays requires parallel, independently-tunable delay lines, which increases the complexity of this analog circuitry. Depending on the antenna configuration and propagation environment, h_{tx} may vary rapidly, requiring frequent retuning of the analog circuitry.

However, more importantly, the discrete number of delay lines limits the reconstruction quality of h_{tx} . Hence, the analog cancellation performance of the Stanford architecture obviously drops in strongly multi-path (i.e., frequency selective) channels [3]. Interestingly, it was shown in [6] that strong passive suppression causes exactly such significantly frequency-selective channel conditions. In other words, the Stanford architecture is not suitable for use after strong passive suppression methods, but is better suited to relatively slow-varying and frequency-flat interference channels, such as the channel created by a circulator, as shown in [3].

Nevertheless, by inspecting (2), we see that it is easy to treat z_s in a second suppression cleanup pass, in the digital domain. This property is the main advantage of the Stanford architecture. In [3] the authors report additional cancellation of -50dB in the second pass of suppression. However, it is not clear how digital suppression can be performed in the presence of a remote signal x , which acts as interference when we are trying to estimate the residual self-interference channel \tilde{h}_s . The presence of the remote signal x is a general problem, affecting any self-interference suppression architecture. However, the presence of x affects the Stanford architecture more severely, since




Fig. 2. Rice architecture, the suppression circuitry is highlighted in gray.

the Stanford architecture relies on achieving high suppression in the second pass in the digital domain, which in turn requires a good estimate of \tilde{h}_s .

Another key point is the scaling with the number of antenna chains in a MIMO system. In this approach, each transmitter-receiver pair requires a dedicated analog circuit, which must be tuned independently. Hence, a total of N^2 such circuits are needed for a system with N antenna chains.

In summary, the Stanford architecture offers high performance through strong second pass of suppression (digital), but it scales poorly with the number of antenna chains, and uses expensive and complex analog circuits whose performance drops in strongly frequency-selective channels.

B. Rice architecture

In Figure 2 we see the structure of the Rice architecture. The received signal is:

$$r_r = \underbrace{(s + e_{tx}) * h_{tx} + (s * \hat{h} + e_{cx}) * h_{cx}}_{\text{self-interference, } z_r} + x * h_x + n + q, \quad (4)$$

where h_{cx} is the impulse response of the wired channel experienced by the cancellation signal, e_{cx} is the transmitter noise of the cancellation chain, and \hat{h} is the digital compensation applied at the cancellation chain, chosen so that $\hat{h} * h_{cx} = -h_{tx}$. Rewriting the self-interference terms:

$$z_r = \underbrace{s * \tilde{h}_r}_{\text{channel estimation error}} + \underbrace{e_{tx} * h_{tx} + e_{cx} * h_{cx}}_{\text{transmitter noise}}, \quad (5)$$

where $\tilde{h}_r = h_{tx} + \hat{h} * h_{cx}$ is the residual channel. The suppression is now given by (7). Since we assume that the transmission and cancellation chains operate at approximately equal power, we have $E[|e_{tx}|^2] \approx E[|e_{cx}|^2]$. We also assume that $\int_{-\infty}^{\infty} |h_{tx}|^2 \approx \int_{-\infty}^{\infty} |h_{cx}|^2$, so that $\int_{-\infty}^{\infty} |\hat{h}|^2 \approx 1$. Hence, (7) becomes:

$$\beta_r \approx \frac{2E[|e_{tx}|^2] \int_{-\infty}^{\infty} |h_{tx}|^2}{E[|s|^2] \int_{-\infty}^{\infty} |h_{tx}|^2} = 2 \times \text{EVM}. \quad (6)$$

Hence, the achievable suppression is limited by the transmitter noise, which is typically characterized in terms of transmitter's error vector magnitude (EVM), defined as $E[|e_{tx}|^2] / E[|s|^2]$.

We show in Section IV that our implementation of the Rice architecture achieves approximately -48dB of active RF suppression. Our implementation uses also significantly simpler circuitry than the analog suppression circuitry presented in [3]. In our implementation only an additional digital to analog converter (DAC) and additional mixer are needed to form the cancellation chain. It is important to note

$$\beta_r = \frac{E \left[\left| s * \tilde{h}_r + e_{tx} * h_{tx} + e_{cx} * h_{cx} \right|^2 \right]}{E \left[\left| s * h_{tx} + e_{tx} * h_{tx} \right|^2 \right]} = \frac{E \left[|s|^2 \right] \int_{-\infty}^{\infty} \left[\tilde{h}_r \right]^2 + \overbrace{E \left[|e_{tx}|^2 \right] \int_{-\infty}^{\infty} |h_{tx}|^2 + E \left[|e_{cx}|^2 \right] \int_{-\infty}^{\infty} |h_{cx}|^2}^{\text{dominant}}}{\underbrace{E \left[|s|^2 \right] \int_{-\infty}^{\infty} |h_{tx}|^2}_{\text{dominant}} + E \left[|e_{tx}|^2 \right] \int_{-\infty}^{\infty} |h_{tx}|^2} \quad (7)$$

that none of the analog components in the Rice architecture need to be tuned. This architecture is in principle not limited in the number of multi-path components that it can cancel, and it can digitally adjust for variations in h_{tx} and h_{cx} .

We also note that in a MIMO system, only one analog circuit is required per receiver chain, because the cancelation signals for all the transmitter chains in the MIMO system can be combined digitally and can be canceled together. Hence, in this architecture, the number of analog components scales with N , rather than N^2 .

However, by looking at (5), we realize that the dominant term in z_r cannot be treated easily in the second pass of suppression, in the digital domain. Hence, the main drawback of the Rice architecture is limited digital suppression. Therefore, in order to make this low-complexity approach practically useful, active RF suppression must push the self-interference to the level of the noise floor.

III. COMPONENTS OF THE SELF-INTERFERENCE SIGNAL

The overall goal of any suppression scheme is to cancel all of the self-interference signal. However, the self-interference signal itself is a sum of several components, with very different power levels. The most important component is the intended transmitted signal, s . Significant components also come from the distortions introduced by the imperfect electronic components in the transmit and receive chains - though all of these are still (linear or non-linear) functions of s . Also, the transmitted signal is modified by the wireless channel as it reaches the receiver. In this section we summarize the components of the self-interference signal, and explain how we suppress each. The same method of reconstruction, for each component, is applied in the first pass (analog RF) and second pass (digital) of suppression.

A. Line of sight (LOS)

The most powerful component of the self-interference signal is likely to come from direct LOS propagation between the transmitter and receiver. This component is the pure transmitted signal s filtered by a single-tap channel h_{tx} . Hence, this component is relatively simple to reproduce, and typically removes a very large portion of the energy in the self-interference signal.

B. Non line of sight (NLOS)

Multipath propagation creates a frequency-selective wireless channel, which has significant energy in NLOS propagation components. After the LOS component has been removed, the NLOS component of the self-interference becomes dominant. The channel is then better represented by an impulse response h_{tx} . Hence, in practice, the LOS component is not reproduced on its own, but instead forms a part of a LOS+NLOS self-interference component given by $s * h_{tx}$. In fact, this component covers all possible linear transformations of the original transmitted signal.¹ This component is also simple to reproduce and removes even more of the energy in the self-interference signal than removing the LOS component only.

¹As long as the delay is shorter than the (finite) length of h_{tx} .

C. Harmonics

Real electronic circuits (e.g., power amplifiers in the transmitter) produce various non-linear components in the transmitted signal. An example of such non-linear components, which was reported in [3], are the harmonics s^n , $n \in \{2, 3, 4, \dots\}$. Each of these components appears at the receiver filtered by its own channel h_n . Therefore, the complete harmonics component of the self-interference signal can then be reproduced as $\sum_{n=2}^{\infty} s^n * h_n$. In all our experiments, we observed that only the harmonics up to $n = 5$ make a noticeable contribution to the self-interference, i.e., are above the noise floor. As noted in [3], only the odd harmonics created on the transmitter side actually appear in the received signal. However, we observed that the even harmonics also appear in the received signal, due to the non-linearities in the receiver.

D. Phase noise

Real hardware oscillators suffer from phase noise, or jitter. A detailed treatment of phase noise and its impact on the performance of FD systems is given in [7], [8]. We note that when the oscillators in the transmitter and canceler chains operate independently, creating each its own carrier from a shared reference signal, they experience two independent phase noise signals. Then, the phase difference between the two oscillators is affected by both noise sources. On the other hand, if both oscillators use the same carrier signal, they are affected by the same phase noise signal. In this case, the phase difference between the two oscillators suffers from much less noise, which only depends on the difference in path length from the carrier source to each of the oscillators, i.e., relative carrier delay. Our hardware platform, described in Section IV-A, allows us to share between the two transmitters either the reference signal from which the carrier signal is generated locally, or the actual carrier signal. The latter of course results in much less phase noise and improved suppression results, thus being a useful and important design guideline for systems using the Rice architecture.

E. Cancelable residual

Once all the above components of the self-interference signal are recreated, and both active RF and digital suppressions are performed, some residual self-interference still remains. We ask the question: how much of that residual is still in principle predictable, and can therefore be canceled? To answer this, we propose an experiment in two parts.

In the first part, the same signal, s , is transmitted repeatedly, for K iterations. No other node is transmitting, so that $x = 0$. The residual, $z^{(k)}$ is recorded on every iteration $k \in \{1, 2, \dots, K\}$. Then, the cancelable component of the residual self-interference is estimated as:

$$\bar{z} = \frac{1}{K} \sum_{k=1}^K z^{(k)}. \quad (8)$$

We confirmed through experiments that \bar{z} changes noticeably only when the transmitted data s changes. Hence, we refer to it as the *constant* component of the residual self-interference. Moreover, the

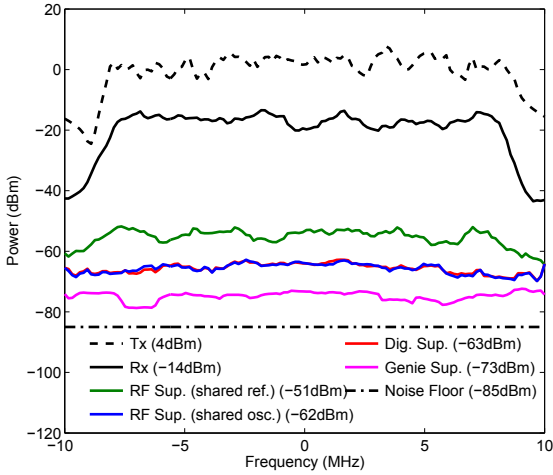


Fig. 3. Successive suppression of self-interference at 4 dBm transmit power and 20 MHz bandwidth using the circulator front-end.

constant component is *uncorrelated* with s^n , $n \in \{1, 2, 3, \dots\}$, meaning that it is not a result of imperfect implementation of the suppression mechanism. Furthermore, we calculate:

$$\tilde{z}^{(k)} = z^{(k)} - \bar{z}, \quad (9)$$

where $\tilde{z}^{(k)}$ is the zero-mean *variable* component of the residual self-interference. From these components we derive the *cancelability metric*:

$$\rho = \frac{\int_{-\infty}^{\infty} |\bar{z}|^2}{\frac{1}{K} \sum_k \int_{-\infty}^{\infty} |\tilde{z}^{(k)}|^2}, \quad (10)$$

which expresses the relative strength of the constant part of the residual self-interference.

A-priori access to \bar{z} , from a genie, would allow us to achieve further suppression. To demonstrate this, in the second part of the experiment the same signal s is transmitted again, another Q times. Now, the genie signal, \bar{z} , which was found in the first part of the experiment is available. It is subtracted from the residual self-interference:

$$\tilde{z}^{(q)} = z^{(q)} - \bar{z} \quad q \in \{1, 2, \dots, Q\}, \quad (11)$$

which achieves further suppression of the self-interference, given by:

$$\beta_g^{(q)} = \frac{\int_{-\infty}^{\infty} |\tilde{z}^{(q)}|^2}{\int_{-\infty}^{\infty} |z^{(q)}|^2}. \quad (12)$$

IV. EXPERIMENTAL RESULTS

A. Hardware platform

Our full-duplex node consists of a National Instruments FlexRIO device with two FlexRIO 5791R RF [9] transceiver modules. The 5791R modules support transmit powers of up to 8 dBm, 100 MHz bandwidth, as well as carrier frequencies from 200 MHz up to 4.4 GHz. The EVM is less than 1.5% RMS (i.e., -36 dB), while the noise power density is -162 dBm/Hz in the 2.4 GHz band where we operate. The first RF module hosts the transmitter and receiver chains, while the second module is used exclusively for the cancellation chain. We have two different antenna front-ends available. The first setup consists of a circulator and a single antenna, which




Fig. 4. Successive suppression of self-interference at 4 dBm transmit power and 20 MHz bandwidth using the antenna front-end with 30 cm antenna separation.

is the same as the setup used [3]. The second consists of two omnidirectional antennae (one for transmission and one for reception) with a separation of 30 cm whose tips are pointing towards each other. This setup is similar to the one used in [4]. Both front-ends include a low noise amplifier (LNA).

The 5791R modules receive and send baseband samples from a computer where all baseband processing is performed in MATLAB. To achieve wideband communication, we use orthogonal frequency-division multiplexing (OFDM). Each OFDM frame contains 20 OFDM symbols with a variable number of carriers, depending on the utilized bandwidth. Specifically, each 10 MHz of bandwidth are split into 256 carriers. We use a carrier frequency of 2.48 GHz. The channel \hat{h} is estimated by piloting. Digital suppression of the s^n , $n \in \{1, 2, \dots, 5\}$, is performed.

B. Suppression results

The suppression results presented in this section are averaged over 20 transmitted OFDM frames.

1) *Circulator front-end*: In Fig. 3 we present indicative spectra of the self-interference signal after cancellation of the different self-interference components when using the circulator based antenna front-end with a transmit power of 4 dBm and a bandwidth of 20 MHz. The legend contains the average suppression achieved after cancellation of the different self-interference components. Based on the specifications of the 5791R, the noise floor when using 20 MHz of bandwidth can be calculated to be -89 dBm. Our measurements of the noise floor indicated an actual noise power of -90 dBm. When the antenna front-ends, which contain an LNA, are attached, the measured noise floor increases to -85 dBm.

We observe that the circulator provides -18 dB of passive suppression, meaning that the interference power at the receiver is -14 dBm. Taking into account the fact that EVM appears twice in (7), we expect to see a floor in the analog suppression at -34 dB. Indeed, in Fig. 3 we see that active RF suppression is able to provide no more than -36 dB of suppression. However, when the carrier is shared between the transmitter and the cancellation chains, thus reducing the phase noise difference between the two chains significantly, the active RF suppression is improved by an additional -12 dB for a total of -48 dB of active RF suppression. This experimental result

$$\beta_{r,n}^{\text{MIMO}} \approx \frac{E[|e_{cx_n}|^2] \int_{-\infty}^{\infty} |h_{cx_n}|^2 + \sum_{k=1}^N E[|e_{tx_k}|^2] \int_{-\infty}^{\infty} |h_{tx_{k,n}}|^2}{\sum_{k=1}^N E[|s_k|^2] \int_{-\infty}^{\infty} |h_{tx_{k,n}}|^2}. \quad (13)$$

confirms the theoretical analysis of the effect of phase noise on active cancellation which is presented in [7], [8]. Moreover, even though the aforementioned analysis concerns narrowband signals, our experimental results indicate that it is also applicable to wideband systems. We observe that, as predicted in [7], [8], digital suppression is not able to reduce the interference signal strength further. This happens because at this point the interference is dominated by the transmitter noise and the accurate estimation of the residual channel \tilde{h}_r is difficult. A more detailed discussion can be found in [7], [8]. The method described in Sec. III-E, which we call *genie suppression*, is able to suppress an additional -10 dB of self-interference. This tells us that there is a strong data-dependent component remaining in the residual self-interference which is, in principle, cancelable provided that it can be modeled appropriately.

We observe that after all of the above suppression methods (except for the hypothetical genie suppression) have been applied, the power of the interference signal is reduced to -63 dBm. Moreover, the genie suppression shows that there is potential to suppress an additional -10 dB, thus reaching -73 dBm. However, this number is still higher than the noise floor of -85 dBm, meaning that the effective noise floor of the receiver is increased significantly.

2) *Antenna front-end*: The antenna front-end provides much better passive suppression, which can potentially reduce the residual interference power to the noise floor, provided that the performance of the active suppression methods is not significantly affected. In Fig. 4 we see that an antenna separation of 30 cm provides passive suppression of approximately -39 dB, which is -21 dB better than the passive suppression achieved by the circulator. Moreover, it is remarkable that active RF suppression can still provide approximately -46 dB of suppression on average, reducing the interference power to -81 dBm. With strong passive suppression, small changes in the environment lead to significant changes in the channel. Thus, even if the estimation is perfect, there is a small mismatch between the real channel and the estimated value because of the time difference between the estimation and the actual usage of the channel state information. This mismatch can be corrected by the digital suppression, which provides an additional -2 dB of cancellation on average. The final interference (plus noise) power is -83 dBm, which is only 2 dB higher than the noise floor. It is significant to note that -85 dB of suppression are achieved *before* the signal reaches the ADC, so the desired remote signal can be sampled using almost the full dynamic range of the ADC.

In Fig. 5 we present the increase in the noise floor for different bandwidths and different transmit powers. We see that, with our specific setup, the Rice architecture is viable since the increase in noise floor is relatively small over a wide range of bandwidths and transmit powers. Moreover, the simplicity and superior scaling behavior of the Rice architecture make it an attractive candidate for MIMO applications.

V. LOW-COMPLEXITY FULL-DUPLEX MIMO

In this section we discuss self-interference suppression in the context of MIMO and we present our low-complexity 2×2 FD-MIMO implementation using the Rice architecture.

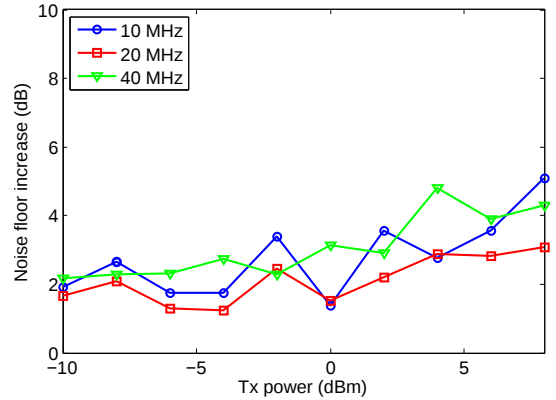


Fig. 5. Noise floor increase for different bandwidths and transmit powers.

A. Self-Interference Suppression for FD-MIMO

In an $N \times N$ FD-MIMO scenario, each transmitter interferes with all receiver chains. Specifically, without cancellation, the received signal at receiver n , $n \in \{1, \dots, N\}$, is (for simplicity, we consider the case where the remote signal x is absent):

$$r_n^{\text{non-sup}} = \sum_{k=1}^N (s_k + e_{tx_k}) * h_{tx_{k,n}} + n_n + q_n, \quad (14)$$

where s_k is the own transmitted signal from antenna k , e_{tx_k} is the transmitter noise of the transmitter chain k , $h_{tx_{k,n}}$ is the wireless channel experienced by the self-interference signal s_k as it travels from transmit antenna k to receive antenna n , n_n is the receiver thermal noise at receiver n , and q_n is the quantization noise introduced in receiver n .

Using the Rice architecture for each receive antenna n , we construct the cancellation signal for each self-interference component separately in the digital domain and then sum these components to obtain the final cancellation signal. As already discussed, by using this approach we can use the simple suppression circuit for each antenna that has already been used in the single antenna scenario. Thus, when suppression is applied, the residual signal at receive antenna n is given by:

$$r_n = \sum_{k=1}^N (s_k + e_{tx_k}) * h_{tx_{k,n}} + \left(\sum_{k=1}^N s_k * \hat{h}_{k,n} + e_{cx_n} \right) * h_{cx_n} + n_n + q_n, \quad (15)$$

where $\hat{h}_{k,n}$ is the digital compensation applied at the cancellation chain of receiver n in order to cancel s_k , and e_{cx_n} is the transmitter noise of the cancellation chain corresponding to receive antenna n .

By extending the analysis in Sec. II-B, one can show that the suppression at receiver n for an $N \times N$ FD-MIMO setup using the Rice architecture is given by (13).




Fig. 6. 2×2 FD-MIMO node schematic, the suppression circuitry for each receiver is highlighted in gray.




Fig. 7. 2×2 FD-MIMO node with the circulator front-end. TX_n , RX_n , and CX_n denote the transmitter, receiver, and cancellation chains for antenna n , respectively.

B. 2×2 FD-MIMO Node

A model and a photograph of our FD-MIMO node are shown in Fig. 6 and Fig. 7, respectively. The cancellation chain for antenna n generates the cancellation signal $\sum_{k=1}^N s_k * \hat{h}_{k,n}$ in the digital domain, which is then upconverted to the RF domain and added to the received signal of each antenna. The analog part of each antenna is the same as the one used in the SISO setup. Thus, the burden of the additional complexity required to suppress N^2 interferers is carried by the digital part of the transceiver, where processing power is cheap, simple, and readily available.

We present suppression results with the circulator based antenna front-end and with the two MIMO antennae separated by a distance of 20 cm. With this setup, we have

$$\int_{-\infty}^{\infty} |h_{tx_{i,j}}|^2 \approx \int_{-\infty}^{\infty} |h_{tx_{k,l}}|^2, \forall i, j, k, l \in \{1, 2\}. \quad (16)$$

Thus, (13) becomes:

$$\beta_{r,n}^{\text{MIMO}} \approx \frac{(N+1) \cdot E[|e_{tx_1}|^2] \int_{-\infty}^{\infty} |h_{tx_{1,1}}|^2}{N \cdot E[|s_1|^2] \int_{-\infty}^{\infty} |h_{tx_{1,1}}|^2} = \frac{3}{2} \times \text{EVM}, \quad (17)$$

which is approximately -1.5 dB better than (7). In Fig. 8, we compare the average suppression achieved by our FD-SISO and our FD-MIMO nodes over different transmit powers when using a bandwidth of 10 MHz. We observe that the average suppression gain of FD-MIMO over FD-SISO is in fact approximately -5 dB, which is even better than what we anticipated.

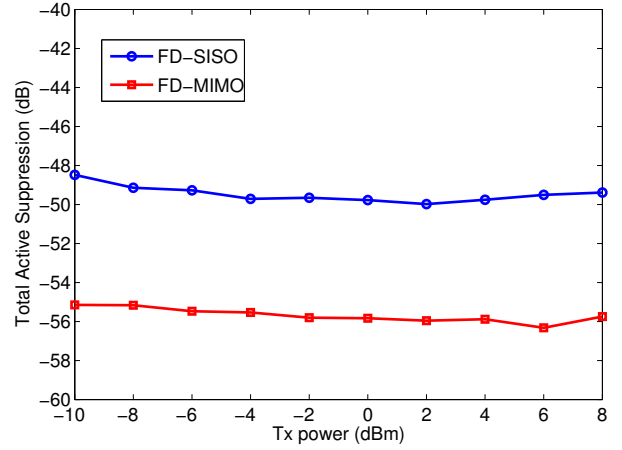


Fig. 8. Comparison of suppression achieved by FD-SISO and 2×2 FD-MIMO at a 10 MHz bandwidth.

VI. CONCLUSIONS

This paper presented our wideband implementation of the Rice architecture, which achieved -48 dB of active RF suppression, and up to -85 dB of total suppression before the ADC. Moreover, we identified a data-dependent deterministic component in the residual interference, which can, in principle, be modeled and canceled, unlocking the potential to improve the currently used cancellation mechanisms. Exploiting the favorable scalability of the Rice architecture, we expanded our implementation to present a low-complexity 2×2 FD-MIMO node. Interestingly, our experimental results show that approximately -5 dB more suppression can be achieved in FD-MIMO than in FD-SISO.

REFERENCES

- [1] J. I. Choi, M. Jain, K. Srinivasan, P. Levis, and S. Katti, "Achieving single channel, full duplex wireless communication," in *Int. Conf. on Mobile Computing and Networking (MobiCom)*, Sep. 2010, pp. 1–12.
- [2] M. Jain, J. I. Choi, T. M. Kim, D. Bharadia, S. Seth, K. Srinivasan, P. Levis, S. Katti, and P. Sinha, "Practical, real-time, full duplex wireless," in *Int. Conf. on Mobile Computing and Networking (MobiCom)*, Sep. 2011, pp. 301–312.
- [3] D. Bharadia, E. McMillin, and S. Katti, "Full duplex radios," in *ACM SIGCOMM'13*, Aug. 2013.
- [4] M. Duarte and A. Sabharwal, "Full-duplex wireless communications using off-the-shelf radios: Feasibility and first results," in *Asilomar Conference on Signals, Systems, and Components*, Nov. 2010, pp. 1558–1562.
- [5] M. Duarte, C. Dick, and A. Sabharwal, "Experiment-driven characterization of full-duplex wireless systems," *IEEE Trans. Wireless Commun.*, vol. 11, no. 12, pp. 4296–4307, 2012.
- [6] E. Everett, A. Sahai, and A. Sabharwal, "Passive self-interference suppression for full-duplex infrastructure nodes," Jan. 2013, submitted to *IEEE Trans. on Wireless Commun.* [Online]. Available: <http://arxiv.org/abs/1302.2185>
- [7] A. Sahai, G. Patel, C. Dick, and A. Sabharwal, "Understanding the impact of phase noise on active cancellation in wireless full-duplex," in *Asilomar Conference on Signals, Systems, and Components*, Nov. 2012, pp. 29–33.
- [8] —, "On the impact of phase noise on active cancellation in wireless full-duplex," 2013, accepted for publication in *IEEE Trans. on Vehicular Techn.* [Online]. Available: <http://arxiv.org/abs/1212.5462>
- [9] National Instruments, *NI 5791R User Manual and Specifications*, May 2013.