Software-based In-Band Full Duplex Relays for IEEE 802.11a/g/p: An Experimental Study

Muhammad Sohaib Amjad and Falko Dressler
School of Electrical Engineering and Computer Science, TU Berlin, Germany
{amjad,dressler}@ccs-labs.org

***Abstract*—Full-Duplex (FD) relay systems have the potential to overcome various system-inherent deficiencies of existing infrastructure relays. In contrast to traditional Half-Duplex (HD) relaying systems, in-band FD relays can simultaneously receive and forward, which improve the channel utilization and reduces the end-to-end delay in a multi-hop network significantly. Despite many theoretical works on FD relaying, there is still no widely available implementation that is compliant with standards such as IEEE 802.11a/g/p. Building upon our previous work, we present such an implementation using the GNU Radio framework. In this paper, we present the performance comparison results with our IEEE 802.11a/g/p compliant relay implementation from a system-level perspective by means of over-the-air experiments in an indoor setup. The first results demonstrate the potential of our Full-Duplex Relay (FDR) implementation and the practical performance in comparison to HD relaying in terms of packet delivery ratio, and physical layer latency.**

I. Introduction

Presently, the growing wireless traffic and high-speed connectivity, are severely bottlenecked by the unpredictable and degrading nature of the wireless channel. The sensitivity of radio signals towards noise, interference, and other time varying channel impairments poses serious challenges for reliable communication – with consequential impact on the overall system performance. The two main criteria that benchmark this reliability check include coverage region and channel capacity; a trade-off is required between these two to ensure error-free wireless communications.

In recent years, multi-hop communication via relay nodes has emerged as an attractive solution to address this channel capacity vs. coverage issue in a band-limited wireless link. Such relay nodes are now even adopted by wireless standards like 3GPP LTE [1] and IEEE WiMAX [2], as they can substantially improve the capacity gains and increase the coverage at the same time. These relays are Half-Duplex (HD) in nature, i.e., they can not simultaneously receive and forward the signal. This results in inefficient channel utilization, and additional resources are required either in time domain (which multiplies the end-to-end latency) or frequency domain (that incurs spectral losses) for interference-free communication.

Over the last decade, in-band Full-Duplex (FD) wireless communication has attracted significant attention, and a considerable amount of research has been done in this domain [3]. The recent advances in signal processing capabilities and antenna technologies have played a pivotal role in enabling full-duplex wireless communications. Different solutions, implementations, and architectures have been proposed, e.g., in [4]–[7], to address the foremost factor impeding in-band FD communications, i.e., the Self-Interference (SI), which intrinsically exist due to radio's own transmission at the same time and frequency. The capacity of a full-duplex radio to transmit and receive simultaneously primarily helps in recuperating with the spectral losses and in improving the effective channel utilization of classical half-duplex systems.

Besides the mentioned advantages, FD communication can also play a critical role in reducing the end-to-end latency of a multi-hop relay network while offering extended coverage and improved capacity gain at the same time. This is essentially helpful in multi-hop latency-critical applications such as road traffic platooning, where a leading vehicle coordinates with its followers, for effective road utilization and fuel-efficient driving [8]. The ability of an in-band Full-Duplex Relay (FDR) to simultaneously receive and forward, greatly reduces the latency of a multi-hop relay system, with least channel utilization, i.e., minimum channel engagement of relaying node. Nevertheless, for optimal performance, the suppression of Loop-back Self-Interference (LSI) (due to simultaneous reception and forwarding) to the receiver's noise floor is a crucial requirement in an FDR system.

This paper extends our previous work [9], where we first modeled and obtained the closed-form expressions for end-to-end Signal-to-Noise Ratio (SNR), and then performed extensive real-time simulations with the proposed IEEE 802.11a/g/p compliant implementation of Amplify and Forward (AF) and Decode and Forward (DF) relaying schemes in both HD and FD modes. In this work, we employ a passive suppression approach and linear digital domain cancellation in the GNU Radio radio framework for real-time system-level evaluation of our standard-compliant implementation via Software Defined Radios (SDRs). To the best of our knowledge, this is the first work that evaluates the performance of FD relaying via real-world experiments using a General Purpose Processor (GPP)-based software implementation adhering to the IEEE 802.11a/g/p standard for WLAN. Our evaluation results demonstrate and underline the potential performance gains offered by FD relays over traditional HD relaying.

Our main contributions can be summarized as follows:

- We improved the LSI suppression module in our preliminary implementation of an FDR system in GNU Radio. The module additionally estimates sampling offset for improved digital domain cancellation of LSI for real-

world experiments with the SDRs, supporting live signal processing and runtime monitoring of LSI cancellation.
- For the first time, we evaluate the performance of an IEEE 802.11a/g/p compliant FDR system, primarily focusing on DF relaying, via real-time SDR-based experiments.
- We investigate the impact of residual LSI on the FDR performance in terms of Packet Delivery Ratio (PDR) and Physical Layer Latency (PLL). We also look into its impact on the noise floor relative to Signal-of-Interest (SoI) and the resultant transmit power requirements from the source node.

II. Related Work

In recent years, multi-hop communication via relay nodes has emerged as an attractive solution to address the ever-increasing demands of higher data rates in wireless networks. In the literature, to overcome the added disadvantages of such relays due to their half-duplex nature, different works have considered approaches such as cooperative decoding for diversity gain [10], and two alternating relay nodes to mimic FD mode [11]. These approaches have aided in improving the spectral usage, however, the spectral losses that Half-Duplex Relay (HDR) incurs are still not entirely compensated.

Recently, the idea of in-band full-duplex relaying has gained significant attention. After all, its implications are qualitatively beneficial in terms of both channel utilization and network latency. In the domain of in-band FD relaying specifically, there is a rather limited amount of literature available on simulative findings, stating the possible improvements and potential gains of FD relaying. In [12], SC-FDE technique based FD relaying is proposed that utilizes AF scheme and is shown to be robust towards LSI via simulative evaluation. Similarly, an adaptive approach for the cancellation of LSI in AF-based relays that do not require the angle-of-arrival information for the temporal filter has been proposed [13]. Likewise, the work in [14], presented a hybrid design for opportunistic FD / HD relaying with transmit power adaptation and showed performance gain over individual mode based systems. In [9], we presented the first complete simulation implementation of IEEE 802.11a/g/p compliant AF and DF based full-duplex relays in the GNU Radio framework. The work is based on an open-source platform and can be studied in all details with the possibility of further modifications for testing and verification of other cancellation algorithms in different channel conditions.

In [15], the first complete FDR design, implementation, and performance evaluation have been presented. The work introduced Construct and Forward (CF) relaying, which unlike the naïve forwarding done by a typical AF relay, forwards the relayed signal such that it constructively adds up with the direct signal (from source) at the destination. Another work [16] introduced a complete prototype implementation for in-band full-duplex relaying, mainly exploiting the LSI techniques, and proposed a novel wave-trap antenna design for passive suppression, and digital domain cancellation stages for LSI suppression to the noise floor. In [17], we presented the

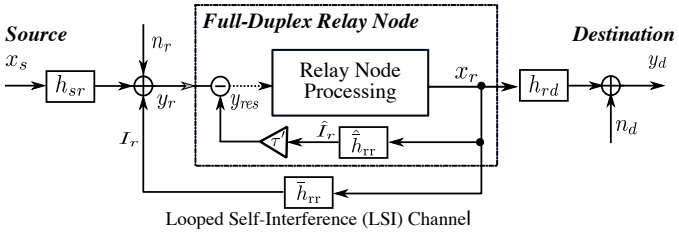

Figure 1. Block diagram of a two-hop transmission link with an FDR node.

first DF-based real-time FDR implementation for an OFDM-based system, considering just QPSK modulation. The system did not yet include the FFT offset correction implementation, which is now incorporated here, in the presented work, for better digital domain LSI cancellation and improved results.

FD relaying solutions, such as [15], are Field-Programmable Gate Array (FPGA)-based and, although these FPGA-based SDRs have deterministic timing and low latency, they have limited flexibility in terms of reprogrammability – besides the expert knowledge needed for handling complex signal processing algorithms in FPGAs. In contrast, the proposed FDR implementation in this work is GPP-based, build upon open-source platform GNU Radio, with the prime advantage of using the same code for both simulations and experiments, therefore, offering a seamless switch between theory and the real-world. Additionally, the framework utilizes high-level programming languages C++ and Python, supports real-time signal processing, a wide variety of SDR radio front-ends, and live visualization of received signals, thus, making it particularly easy to use, modify, and debug.

III. Loop-back Self-Interference Modeling

We begin by revisiting our previous model [9] of a two-hop relay system, which includes a source node communicating with a destination node via an intermediary relay node, as shown in Figure 1. The received signal samples (y_r) at the input of the relay node include the SoI samples (x_s) from source node and the LSI samples (I_r) of the ongoing/forwarding transmission at the FDR node, and can be written as

$$y_{\mathrm{r}}[i] = x_s[i] * \bar{h}_{sr} + n_r[i] + I_r[i - \tau']. \quad (1)$$

In Equation (1), τ' models the delay incurred by the front-ends, the channel $(\bar{h}_{rr})$ and the decoding process (DF scheme) at the FDR node, $\bar{h}_{sr}$ are the source–relay channel coefficients, and n_r are the Additive White Gaussian Noise (AWGN) noise samples. The residual signal (y_{res}) after LSI suppression can then be obtained as

$$y_{\mathrm{res}}[i] = x_s[i] * \bar{h}_{sr} + n_r[i] + I_r[i - \tau'] - \hat{I}_r[i - \tau'], \quad (2)$$

where the estimated LSI samples $(\hat{I}_r[i])$ are generated as

$$\hat{I}_r[i] = x_r[i] * \hat{\bar{h}}_{rr}. \quad (3)$$

Here, x_r are the samples to be forwarded by the FDR node, and $\hat{\bar{h}}_{rr}$ are the estimated coefficients of the LSI channel. Since the samples to be forwarded $x_r[i]$ are readily available in the digital domain for the regeneration of LSI $(\hat{I}_r)$, therefore,

buffering (τ') of the reconstructed samples ($\hat{I}_r$) is required to compensate for the delay experienced by the received LSI (I_r). Note that, if τ' in (2) is not acquired correctly, then the subtraction of non-synchronized estimated LSI $\hat{I}_r$ from the received LSI I_r can drive the system towards instability. This synchronization of the estimated and received LSI is mostly overlooked (i.e., implicitly assumed to be synchronized) in the available simulation and analytical studies. For real-time systems however, it is the fundamental requirement in FDR implementations; more details are provided in Section IV.

A. Loop-back Self-Interference (LSI) Regeneration

The regeneration of LSI ($\hat{I}_r[i]$) is a two step process, and in this work it is performed in the time-domain. In the first step, an estimate of the LSI channel ($\hat{\bar{h}}_{rr}$) is obtained. Whereas, in the second step, the known samples $x_r[i]$ are filtered in time-domain with the estimated LSI channel coefficients $\hat{\bar{h}}_{rr}$ as shown in Equation (3). The second step enforces channel impairment effects on the forwarded samples x_r in order to generate approximate LSI ($\hat{I}_r$). Thus, the reconstructed self-interference samples $\hat{I}_r[i]$, essentially carry similar channel properties as that held by the received LSI samples $I_r[i]$

1) LSI Channel Estimation: For the estimation of LSI channel, we have employed the Least Squares (LS)-based time-domain estimation technique. The LS approach basically acquires the Channel Impulse Response (CIR) estimate $\hat{\bar{h}}_{rr}$ through the Long Training Sequence (LTS) symbol embedded in the WLAN frame structure. For the in-band full-duplex communication, this estimation process is typically completed during the training transmissions, i.e., no emissions from the other radio or in-case of FD relaying no transmissions from the source node. From Figure 1, the received samples $y_r[i]$ during training transmissions, i.e., $x_s[i] = 0$, can be written as

$$y_r[i] = x_r[i] * \bar{h}_{rr} + n_r[i], \tag{4}$$

i.e., only looped-back self-interference samples ($I_r[i]$) are present. Using (4), the received LTS symbol samples can be presented as

$$\bar{y}^l_{r\text{-}lts} = \bar{x}^l_{r\text{-}lts} * \bar{h}^P_{rr} + \bar{n}^l_{r\text{-}lts}, \tag{5}$$

where l indicates the number of LTS samples (equal to FFT size, i.e., 64 in this case), and P represents the length of channel taps to be estimated, which typically corresponds to selected Cyclic Prefix (CP) length. For predetermined and known (standardized) $\bar{x}^l_{r\text{-}lts}$ samples, the time-domain convolution in (5) can be transformed into matrix multiplication as

$$\bar{y}^l_{r\text{-}lts} = \mathbf{X}^{l\times P} \cdot \bar{h}^P_{rr} + \bar{n}^l_{r\text{-}lts}. \tag{6}$$

In Equation (6), $\mathbf{X}^{l\times P}$ is the Toeplitz matrix of order $l \times P$, constructed (see [6]) using the known sent LTS samples. The time-domain LS channel estimate is thus acquired as

$$\hat{\bar{h}}^P_{rr} = \mathbf{X}^{l\times P^\dagger} \cdot \bar{y}^l_{r\text{-}lts}, \tag{7}$$

where $\mathbf{X}^{l\times P^\dagger}$ is the Moore-Penrose (pseudo) inverse of $\mathbf{X}^{l\times P}$ and $\bar{y}^l_{r\text{-}lts}$ are the received LTS samples. Notice that the sent LTS samples ($\bar{x}^{lts}_l$) are fixed and already known, thus the matrix $\mathbf{X}^{l\times P^\dagger}$ can be precomputed and stored in advance, without additional computational requirement. Additionally, the error magnitude between the estimated and actual LSI channels is obtained as

$$\|\bar{e}_{rr}\|^2 = \|\hat{\bar{h}}^P_{rr} - \bar{h}^P_{rr}\|^2 = \|\mathbf{X}^{l\times P^\dagger} \cdot \bar{n}^l_{r\text{-}lts}\|^2, \tag{8}$$

i.e., the main source of error in the estimation of LSI channel ($\bar{h}_{rr}$) is the receiver noise values.

IV. System Design and Implementation

For the experimental evaluation of Time Division Duplex (TDD)-based HDR and in-band FD relays, we implemented the DF relaying scheme in the GNU Radio-based signal processing framework. We choose this platform because of its wide-spread use in the community, and its capacity to do rapid prototyping – meanwhile supporting both simulation study as well as experimental testing using SDR.

For the baseband modulation/demodulation, and the DF scheme at the relay node, we used the GNU Radio-based open-source stack for IEEE 802.11a/g/p WLAN developed by Bloessl et al. [18], as discussed in [9]. The key parameters within this transceiver design are listed in Table I. For an FDR node, the suppression of LSI is required for which we have implemented a core block in GNU Radio, and modified it further, for better LSI cancellation in the digital domain.

A. LSI Cancellation Block

Figure 2 shows our time-domain based LSI cancellation block that we adapted to work in the scope of the GNU Radio framework. It basically estimates the LSI channel and takes the known relay forwarded samples to reconstruct approximate LSI (only linear component). The block first transmits $C_k + 1$ training packets at the relaying node (during training transmissions) for the estimation of SI channel, and for stabilizing the sub-blocks such as signal synchronizer. In the figure, C_k represents the number of training packets and k indicates the frequency of the LSI channel estimation process. During the transmission of training packets, emissions from the source node are turned off until the relay switches to full-duplex relaying mode.

1) Preprocessing: The preprocessing block performs the tasks which utilize the known parameters for information extraction, and are independent of real-time processing. It first computes the Inverse Fast Fourier Transform (IFFT) of known LTS symbols enclosed in the WLAN packet structure,

Table I
KEY PARAMETERS OF OPEN-SOURCE STACK FOR IEEE 802.11A/G/P.

Parameters	
Modulations	B-PSK, Q-PSK, 16-QAM & 64-QAM
Code Rates	1/2, 3/4, 2/3
Sampling Frequency [MHz]	10, 20
Data Rates [Mbit/s]	6, 9, 12, 18, 24, 36, 48, 54
Cyclic Prefix Length	16 samples
PLCP (Preamble + Header)	(4 + 1) OFDM Symbols

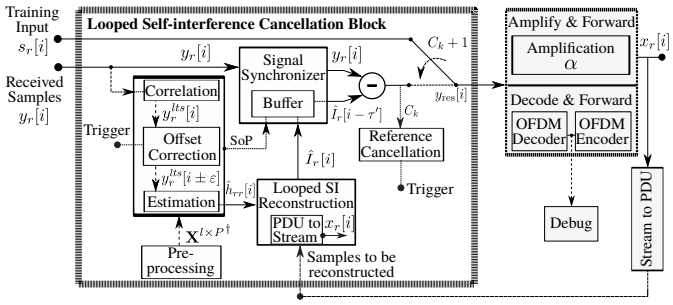

Figure 2. Baseband level block diagram of the implemented novel LSI cancellation module for FDR implementation.

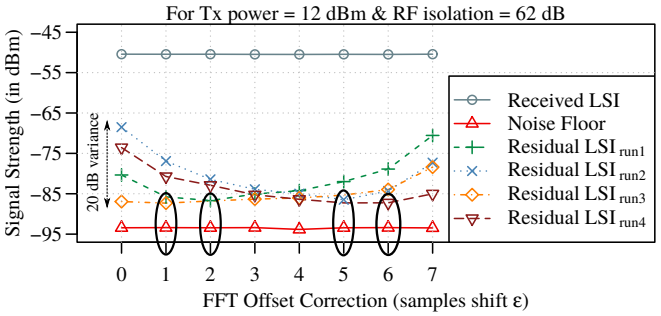

Figure 3. Impact of FFT offset correction on the residual signal. For these measurements, Ettus B210 USRP SDR has been used. The results are obtained for fixed 12 dBm transmit power, and for an RF isolation ≈ 62 dB. The received signal (around −50 dBm) for each run, can be suppressed close to −87 dBm, roughly 7 dB above the receiver's noise floor (−94 dBm) for certain offset value ε. Without this offset correction (see results in [19]), the residual signal strength can go as high as −67 dBm, which is roughly 20 dB more than the maximum achievable digital domain suppression values.

hence converting it into time-domain samples. Afterward, the obtained samples are used to form a Toeplitz matrix $\mathbf{X}^{lxP}$, and finally to compute the Moore-Penrose (pseudo) inverse $\mathbf{X}^{l\times P\dagger}$ of the Toeplitz matrix. Here, l is the same as IFFT size and P is set as half of CP, the values of both are listed in Table I.

2) Estimation: The estimation block operates only during the training transmissions, i.e., C_k packets transmission. It starts with the correlation of received samples $y_r[i]$ with the known LTS samples $\bar{x}_l^{lts}[i]$ to determine the Start-of-Packet (SoP). Once SoP is determined, it then extracts the received LTS samples $y_r^{lts}[i]$, which later undergoes the process of iterative FFT offset correction. The amount of FFT offset is calculated through the reference cancellation block in Figure 2, which holds the offset value providing maximum cancellation as reference. In each iteration (for offset correction) it checks the averaged achieved cancellation, compares it with the reference cancellation value, and overrides the previous cancellation and FFT offset values $(y_r^{lts}[i \pm \varepsilon])$ if the current one is offering better results. The typical offset (ε) value goes as high as 8 samples. Figure 3 shows the impact of offset correction on the residual signal strength. This offset calibration is only required to be done once, and it is finished during the training transmissions. Afterwards, the estimated offset does not change, and thus, the employed iterative calibration does not contributes towards additional overhead.

3) LSI Reconstruction: The reconstruction block as a first step translates the message containing relay forwarded samples $x_r[i]$ into streaming samples and then filters them with the computed LSI channel estimate $\hat{h}_{rr}[i]$ to reproduce approximate LSI samples $\hat{I}_r[i]$.

4) Signal Synchronizer: The signal synchronizer block basically compensates for any delay τ' between the reconstructed LSI samples $\hat{I}_r[i]$, and the received samples $y_r[i]$ (which in-fact are the received LSI samples), during the training transmissions. The fed back samples $x_r[i]$, necessarily arrive earlier compared to the received LSI samples, which are typically delayed by the Tx – Rx front-ends, and software to hardware translation of samples. For this reason, the synchronizer starts buffering the reconstructed samples $\hat{I}_r[i]$ and waits for an SoP indicator to release them. It is worth mentioning here that the synchronizer block calculates the required buffer length during the training session, i.e., no transmissions from source, and once the buffer length is determined it does not change because the delay from relay Tx – Rx remains the same.

After synchronization, the reconstructed relay forwarded samples $\hat{I}_r[i-\tau']$ are subtracted from the received samples and pushed to the relaying block; given that the training period is finished. The relaying block implements both AF and DF schemes, and a debugger is also there to acknowledge both the packet reception and decoding at the relay node.

B. Passive Suppression

Since both the receiving and forwarding chains are overly close at the relay node, the LSI signal is considerably stronger then the SoI arriving from a distant source, and if not repressed to an extent, it can consume the entire dynamic range of the Analog-to-Digital Converters (ADCs) in the received signal processing path. Therefore, the passive suppression stage is rather pivotal.

In this work, to effectuate the passive suppression of the direct/leaked LSI component at the relay node of our FDR system, we have employed the basic antenna isolation approach, which provides a passive suppression of approx. 62 dB. We placed an aluminum foil wrapped Balsa foam between the receiving and forwarding dipole antennas for the isolation of on-board Tx–Rx front-ends within a Universal Software Radio Peripheral (USRP) as shown in Figure 5. It is important to mention here that the employed passive suppression technique in this work provides the lower bound of the achievable performance, and can be very easily replaced with high-end and sophisticated passive suppression and active analog cancellation approaches such as [20], and it is certainly not the main focus of this work.

V. Practical Real-time Performance

For practical performance evaluation of our IEEE 802.11a/g/p compliant FDR implementation, we have conducted detailed experiments in an indoor environment. In the experiments, we used three USRP B210 SDRs as source, relay, and destination nodes, as illustrated in Figure 4. The performance of our FDR implementation is evaluated in Nonline-of-Sight (NLOS) experimental settings. In each iteration, from

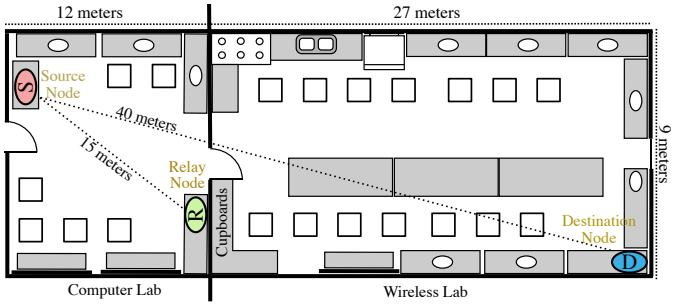

Figure 4. Floor plan for the indoor NLOS experimental settings.

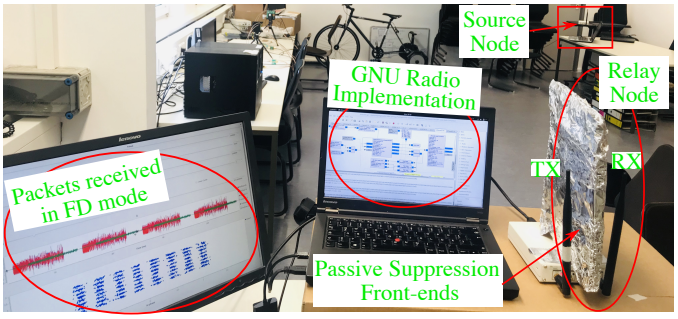

Figure 5. Snapshot of our working hardware setup.

the source node, we transmitted 100 packets per transmission, where each packet comprises 250 B payload, 3 B header, 4 B CRC, and 4 synchronization symbols. The transmissions were repeated 10 times for each Modulation and Coding Scheme (MCS) and with different transmit power levels. All relevant parameters are listed in Table II.

We have only evaluated the Decode and Forward (DF) scheme, because of the Amplify and Forward (AF) scheme's poor analytical and simulative performance as demonstrated in [9]. Here, it is worth mentioning that since decoding delay in DF relaying scheme is the same regardless of HD or FD transmission mode, therefore, its impact is not studied in this work. Figure 5 shows our working hardware setup, indicating the transmitting source node, the receiving and forwarding relay node, the running GNU Radio implementation, and the real-time SoI in green (after LSI suppression) with the resultant decoded constellation for 64-QAM $3/4$. It is worth pointing out here that in the real-time experiments, to avoid interference from WiFi devices and other wireless systems operating in the ISM band, we have used the 868 MHz frequency band, instead of the defined 2.4 GHz and 5 GHz bands in the standard.

For the estimation of the LSI channel, and for FFT offset

Table II
HARDWARE-SPECIFIC PARAMETERS FOR THE NLOS EXPERIMENTS

Carrier Frequency	868 MHz
Sampling Frequency	2 MHz (limited CPU processing power)
Receiver's Noise Floor	−94 dBm
S–R Distance	15 m
S–D Distance	40 m
RF Isolation	approx. 62 dB
Digital LSI Suppression	up-to 39 dB

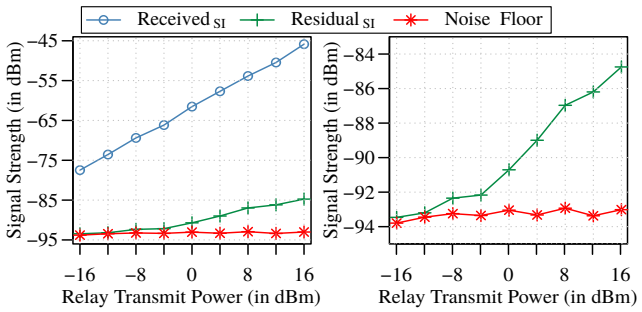

Figure 6. The LSI suppression performance of the presented FDR implementation, in the digital domain for increasing levels of transmit power at the relay node.

correction, we sent 20 training packets (i.e., $C_k = 20$ transmissions). Like any typical FD system, the (short) training packets containing only training sequence symbols are used for the estimation here. How often such training is required highly depends on a fully functional MAC protocol, and it basically defines the additional overhead due to training. These overheads, nevertheless, are quite small. For instance, a standard IEEE 802.11 compliant training packet with only training symbols contains two Short Training Sequence (STS) and two LTS symbols, i.e., 80 x 4 x 20 = 6400 samples are used. Now, assuming that the training is required every 10 ms for a 20 MHz standard sample rate. Then, the total overhead due to training symbols is only 3.2 %. If the training is required every 100 ms, then the overhead is only 0.32 %. As mentioned before, the exact calculations depend on the MAC protocol and lie outside the scope of this work.

A. LSI Suppression at Relay Node

Figure 6 presents the LSI cancellation acquired in the digital domain with different relay node transmit power levels. The computed noise floor of B-210 USRP SDRs at a sampling frequency of 2 MHz is close to −94 dBm, i.e., the red plot. In the figure, it can be observed that for low transmit power levels (up to −12 dBm), the received LSI is suppressed to the receiver's noise floor. However, a non-linear increase in the residual LSI is observed at higher power transmissions due to the following reasons.

First, the non-linear factor introduced by the transmit chain amplifier at higher gain values becomes more significant, resulting in increased strength of the non-linear LSI component, which is not modeled in the current implementation. Secondly, the acquired RF isolation of 62 dB is not sufficient enough (reported as high as 73 dB in the literature), and with higher power transmissions inadequate isolation gets more evident. By further integrating the non-linear LSI component modeling, and employing sophisticated Radio Frequency (RF) isolation approaches such as dual-port dual polarized slot coupled antenna, this residual LSI can be eliminated even for higher relay node transmit power level.

B. Packet Delivery Ratio

Figure 7 demonstrates the PDR (for each MCS) obtained at the relay node operating in both FD and HD modes. The

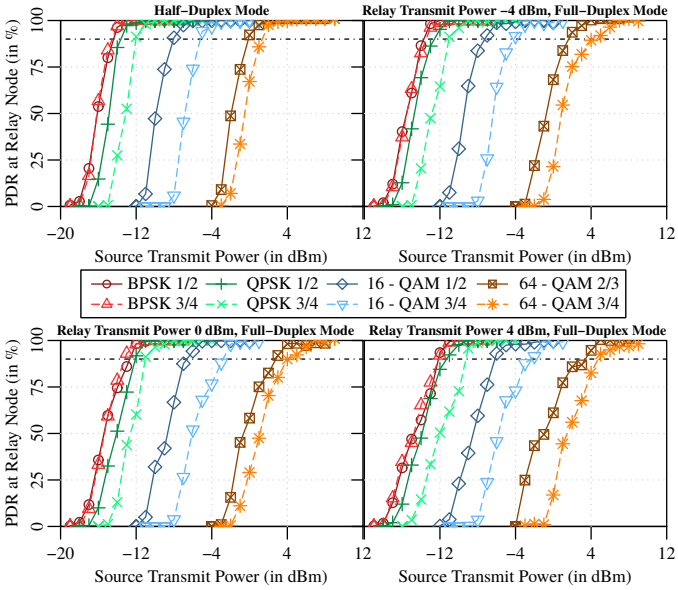

Figure 7. Experimentally measured PDR performances for each MCS with DF-based IEEE 802.11 a/g/p compliant HDR and FDR implementations at different relay node transmit power levels.

titles of the four plots list the operational mode, and the relay node transmit power level (more relevant for FD case). In the plots, PDR 100 % indicates that all packets have been correctly detected and decoded. The dashed horizontal line marks the 90 % PDR level. Also, for clarity in the plots, the 95 % confidence intervals are not shown.

In the figure, it can be seen that the PDR plots with HD and −4 dBm FD modes are relatively similar. There is a difference of roughly 1 dB in the performance of each MCS certainly because of the non-negligible residual LSI. Also, the PDR performances of each MCS at 0 dBm and 4 dBm FD modes are slightly degraded. In both cases, 100 % PDR is achieved at relatively higher source transmit power levels. The reason for this depreciation in the performance is that when more transmit power is used at the relay node, the intensified residual LSI raises the noise-plus-interference level for the SoI arriving from the source node. Hence, higher power is needed from the source node to overcome the raised interference level and maintain 100 % PDR.

Furthermore, if we look into the relative sensitivity levels obtained for each MCS by choosing BPSK 1/2 as a reference scheme, it can be seen that these relative sensitivity levels meet the levels provided in the 802.11[1]. For example, considering the sensitivity value provided in the *802-11ac-receiver-minimum-input-sensitivity-test* for BPSK 1/2 as a reference, the additional gain required by 64-QAM 3/4 is (−65 dBm − (−82 dBm) = 17 dB). Now, in our results (Figure 7), it can be seen that roughly 17 dB more power is required from the source node for 64-QAM 3/4 to achieve similar PDR performance as that obtained with BPSK 1/2, regardless of the residual LSI that is effecting all the modulation and coding

[1]https://uk.mathworks.com/help/wlan/examples/802-11ac-receiver-minimum-input-sensitivity-test.html

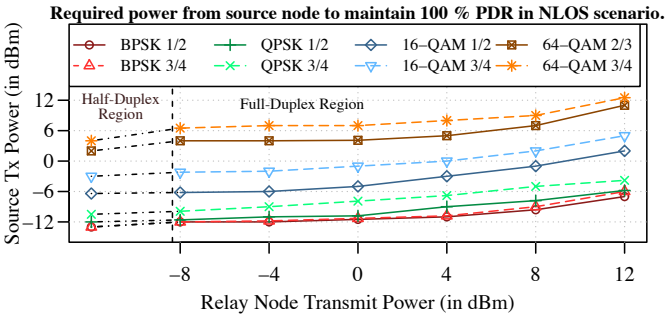

Figure 8. Required source transmit power levels to maintain 100 % PDR with each MCS at the relay node for a given relay transmit power. The half-duplex PDR performance is independent of relay transmit power levels due to time separate transmissions and receptions.

schemes rather equally. It is worth mentioning here that the exact sensitivity levels depend on the receiver's noise floor, which is influenced by the receiver's noise figure and SDR front-end calibration, it is hard to obtain precise minimum sensitivity. Nevertheless, the relative sensitivity values meet the 802.11 standard requirements.

C. Impact of Residual LSI on PDR

Figure 8 shows the required transmit power to retain a 100 % PDR with each MCS at the relay node. The vertical dashed line in the plot separates the FD mode performance with HD mode performance. Ideally, with complete LSI elimination, the source transmit power requirement in the FD case should have been the same as in HD mode. Nevertheless, the non-linearly growing demand for the source power (in case of FD mode) is due to the increasing levels of residual LSI at higher relay node transmit power. This as a consequence, increases the noise-plus-interference level for the SoI, and a higher transmit power is required from the source to maintain the 100 % PDR.

We would like to point out here that the presented system utilizes a very simple RF isolation approach and no cancellation in the analog domain is implemented. The results presented here essentially show the lower bound of the system's performance. Also, the required source transmit power to maintain 100 % PDR at the relay node directly relates to the residual LSI indicated in Figure 6. Since the implementation is based on an open-source framework, further research with more sophisticated suppression approaches can easily be built upon the current implementation to optimize the system's performance.

D. Physical Layer Latency

Figure 9 outlines the minimum achievable PLL comparison with HD and FD relaying in our experimental measurements. It is important to mention here that the delay due to the software-to-hardware translation of samples, the Tx–Rx front-ends, and the propagation, is the same regardless of HDR and FDR implementation, therefore, they are not considered here. From the figure, it can clearly be seen that the PLL decreases with higher transmit power from the source node. The PLL depends on the packet size, or more precisely the number of samples

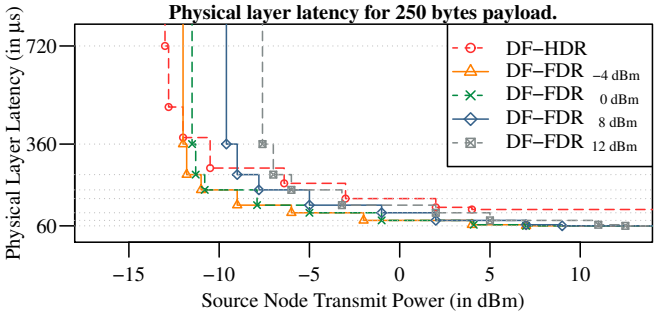

Figure 9. Achievable physical layer latency over the two-hop relaying system in HDR and FDR configurations, with different relay node transmit power levels in the FDR case.

sent per packet, and the decoding delay of the DF scheme. The hardware delay (in the order of a few samples) is not evaluated here. With a higher Tx power resulting in a better received SNR, higher-order modulation-coding schemes defined in the IEEE 802.11 standard (from BPSK 1/2 to 64-QAM 3/4) can be adopted. This means the same number of bytes (the original packet) can be encoded in fewer symbols and, therefore, less samples to transmit (i.e., a reduced transmission time). This reduces the physical layer latency, but at the expense of more transit power.

Additionally, It can be gathered from the figure that although HDR starts receiving the packet roughly 1 dB before FDR, i.e., −12.1 dBm, however, the minimum PLL drops down much faster with FDR. Even at higher relay node transmit power with most residual LSI magnitude, take DF-FDR$_{12\ \mathrm{dBm}}$ as an example, where for the same source transmit power (−6 dBm and above), FDR has lower PLL as compared to HDR. These results conclusively highlight the possible gains with full-duplex relaying in-terms of lower end-to-end latency, which further improves with better LSI suppression.

VI. Conclusion

We experimentally evaluated the performance of IEEE 802.11a/g compliant software-based Full-Duplex Relay (FDR) systems, in particular focusing on Decode and Forward (DF) relaying. Contrary to the earlier works, we further realized both real-time simulation and Software Defined Radio (SDR)-based experimental results with our GNU Radio implementation of the Half-Duplex (HD) and Full-Duplex (FD) relaying systems. Even though there is a need for more advanced RF isolation between the two antennas, our system is able to achieve very good performance results in the indoor setup. Our results show the potential gains offered by FD relays in-terms of efficient channel utilization and reduced physical layer latency from the system-level's perspective. In particular, given that the LSI is sufficiently suppressed, DF-based FDR outperformed the HD relaying systems.

The employed GNU Radio radio framework supports real-time signal processing, and the same code can be used, for both simulations and real-world experiments. The presented IEEE 802.11 compliant FD relaying implementation can also be evaluated in full capacity within the real-time simulation mode, under different channel conditions and impairment sources. Thus, different algorithms to deal with changing channel conditions can be tested in full detail.

References

[1] "Further advancements for E-UTRA physical layer aspects," 3rd Generation Partnership Project, TR 36.814 V9.0.0 (2010-03), Mar. 2010.

[2] "Local and metropolitan area networks," IEEE, Standard P802.16j-06/026r4, May 2007.

[3] Z. Zhang, K. Long, A. V. Vasilakos, and L. Hanzo, "Full-duplex wireless communications: Challenges, solutions, and future research directions," *Proceedings of the IEEE*, vol. 104, no. 7, pp. 1369–1409, 2016.

[4] D. Bharadia, E. McMilin, and S. Katti, "Full Duplex Radios," *ACM SIGCOMM Computer Communication Review (CCR)*, vol. 43, no. 4, pp. 375–386, Oct. 2013.

[5] E. Ahmed and A. M. Eltawil, "All-digital self-interference cancellation technique for full-duplex systems," *IEEE Transactions on Wireless Communications (TWC)*, vol. 14, no. 7, pp. 3519–3532, Jul. 2015.

[6] M. S. Amjad, H. Nawaz, K. Ozsoy, O. Gurbuz, and I. Tekin, "A Low-Complexity Full-Duplex Radio Implementation With a Single Antenna," *IEEE Transactions on Vehicular Technology (TVT)*, vol. 67, no. 3, pp. 2206–2218, Mar. 2018.

[7] J. Zhou, N. Reiskarimian, J. Diakonikolas, T. Dinc, T. Chen, G. Zussman, and H. Krishnaswamy, "Integrated Full Duplex Radios," *IEEE Communications Magazine (COMMAG)*, vol. 55, no. 4, pp. 142–151, 2017.

[8] F. Dressler, F. Klingler, M. Segata, and R. Lo Cigno, "Cooperative Driving and the Tactile Internet," *Proceedings of the IEEE*, vol. 107, no. 2, pp. 436–446, Feb. 2019.

[9] M. S. Amjad and F. Dressler, "Performance Insights on IEEE 802.11a/g Compliant Software-based In-Band Full Duplex Relay Systems," in *IEEE Wireless Communications and Networking Conference (WCNC 2019)*. Marrakesh, Morocco: IEEE, Apr. 2019.

[10] J. Laneman, D. Tse, and G. Wornell, "Cooperative Diversity in Wireless Networks: Efficient Protocols and Outage Behavior," *IEEE Transactions on Information Theory*, vol. 50, no. 12, pp. 3062–3080, Dec. 2004.

[11] B. Rankov and A. Wittneben, "Spectral efficient protocols for half-duplex fading relay channels," *IEEE Journal on Selected Areas in Communications (JSAC)*, vol. 25, no. 2, pp. 379–389, Feb. 2007.

[12] Y. Liu, Y. Yue, X.-G. Xia, C. Kang, and H. Zhang, "SC-FDE based full-duplex relay communication robust to residual loop interference," *IEEE Wireless Communications Letters*, pp. 538–541, 2017.

[13] E. Antonio-Rodríguez and R. López-Valcarce, "Cancelling self-interference in full-duplex relays without angle-of-arrival information," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP 2013)*. Vancouver, Canada: IEEE, May 2013, pp. 4731–4735.

[14] T. Riihonen, S. Werner, and R. Wichman, "Hybrid full-duplex/half-duplex relaying with transmit power adaptation," *IEEE Transactions on Wireless Communications (TWC)*, vol. 10, no. 9, pp. 3074–3085, 2011.

[15] D. Bharadia and S. Katti, "Fastforward: Fast and constructive full duplex relays," *ACM SIGCOMM Computer Communication Review (CCR)*, vol. 44, no. 4, pp. 199–210, 2015.

[16] D. Korpi, M. Heino, C. Icheln, K. Haneda, and M. Valkama, "Compact inband full-duplex relays with beyond 100 dB self-interference suppression: Enabling techniques and field measurements," *IEEE Transactions on Antennas and Propagation*, vol. 65, no. 2, pp. 960–965, 2016.

[17] M. S. Amjad and F. Dressler, "Software-based Real-time Full-duplex Relaying: An Experimental Study," *IEEE Transactions on Green Communications and Networking (TGCN)*, 2020, online: 10.1109/TGCN.2020.2966140.

[18] B. Bloessl, M. Segata, C. Sommer, and F. Dressler, "An IEEE 802.11a/g/p OFDM Receiver for GNU Radio," in *ACM SIGCOMM 2013, 2nd ACM Software Radio Implementation Forum (SRIF 2013)*. Hong Kong, China: ACM, Aug. 2013, pp. 9–16.

[19] M. S. Amjad and F. Dressler, "Experimental Insights on Software-based Real-Time SI Cancellation for In-Band Full Duplex DF Relays," in *IEEE International Conference on Communications (ICC 2019)*. Shanghai, China: IEEE, May 2019.

[20] K. E. Kolodziej, B. T. Perry, and J. S. Herd, "In-band full-duplex technology: Techniques and systems survey," *IEEE Transactions on Microwave Theory and Techniques*, vol. 67, no. 7, pp. 3025–3041, 2019.