SCATTER PHY: A Physical Layer for the DARPA Spectrum Collaboration Challenge

Felipe A. P. de Figueiredo*§, Dragoslav Stojadinovic†§, Prasanthi Maddala†§, Ruben Mennes‡, Irfan Jabandžić*, Xianjun Jiao*, and Ingrid Moerman*

*Ghent University - imec, IDLab, Department of Information Technology, Ghent, Belgium.

†WINLAB, Rutgers University, New Brunswick, NJ, USA.

‡Department of Mathematics and Computer Science, University of Antwerp - iMinds, Antwerp, Belgium

Email: *[felipe.pereira, irfan.jabandzic, xianjun.jiao, ingrid.moerman]@ugent.be, †[stojadin, prasanti]@winlab.rutgers.edu, ‡ruben.mennes@uantwerpen.be

§These authors contributed equally to this work.

***Abstract*—DARPA has started the Spectrum Collaboration Challenge with the aim to encourage research and development of coexistence and collaboration techniques of heterogeneous networks in the same wireless spectrum bands. Team SCATTER has participated in the challenge from its beginning and is currently preparing for the final phase of the competition. SCATTER's physical layer (SCATTER PHY) has been developed as a standalone application, with the ability to communicate with higher layers of SCATTER's system via ZeroMQ, and uses USRP X310 software-defined radio devices to send and receive wireless signals. SCATTER PHY relies on USRP's ability to schedule timed commands, uses both physical interfaces of the radio devices, utilizes the radio's internal FPGA board to implement custom high-performance filtering blocks in order to increase its spectral efficiency as well as enable reliable usage of neighboring spectrum bands. This paper describes the design and main features of SCATTER PHY and showcases the experiments performed to verify the achieved benefits.**

***Index Terms*—Cognitive Radios; Spectrum Sharing; Coexistence; Experimental Evaluation.**

I. Introduction

In the era of a growing number of wireless communication networks and protocols, spectrum scarcity represents an ever-increasing challenge in the research community. DARPA's Spectrum Collaboration Challenge (SC2) motivated teams from the entire world to build networks able to not only coexist with others in the same spectrum but also to use collaboration to ensure that each network maximizes its throughput while facilitating spectrum usage and ensuring its availability when necessary to other networks and systems.

The SCATTER team has designed and built a multi-layered system, with physical, MAC and other layers being completely autonomous but interconnected with a publish-subscribe messaging system, in particular ZeroMQ. Thus, SCATTER's physical layer (PHY) was designed as an independent system, able to transmit and receive data packets over the air using USRP software-defined radio (SDR) devices. The design of SCATTER PHY uses the existing srsLTE software library as the basis for its communication, and builds on top of it with various features and modifications targeted for the SC2, such as the ability to communicate with higher layers of the system (MAC, AI, etc.), and using custom-built FPGA filters and other blocks to reduce out-of-band emissions and maximize channel utilization. This paper describes the design and features of SCATTER PHY and experimentally verifies the benefits achieved with the outlined design decisions.

The remainder of the paper is organized as follows: section II describes the design of SCATTER PHY in detail and specifies its main features. Section III showcases the experimental verification of the achieved benefits. Section IV concludes the paper.

II. SCATTER PHY

The high-level architecture of the SCATTER PHY is depicted in Figure 1. The SCATTER PHY is implemented using Universal Software Radio Peripheral (USRP) Hardware Driver (UHD) software Application Programming Interface (API) [15] and runs on top of Ettus USRP X family of software-defined radio (SDR) devices including NI's RIO platforms [3, 4] and communicates with it through the UHD driver and its APIs [9]. As can be seen in the figure, the individual PHY modules are connected to the ZeroMQ (Data/Control) module, also known as 0MQ, which interconnects the SCATTER PHY with the MAC layer through the ZeroMQ bus [2]. This module manages the exchange of control and statistics messages between the SCATTER PHY and MAC layer.

Communication with the SCATTER PHY is entirely implemented through a well-defined interface designed with Google's Protocol Buffers (protobuf) [1] for data serialization coupled with the ZeroMQ messaging library [2] for distributed exchange of control, statistics and data messages. Implementing the ZeroMQ push-pull pattern allows local or remote MAC layer's real-time configuration of several parameters and reading of several pieces of information/statistics provided by the SCATTER PHY. Based on the ZeroMQ logic, PHY and MAC layers are able to exchange control and data messages following a non-blocking communication paradigm. The SCATTER PHY was designed to be completely decoupled and independent of the MAC layer module, not posing any constraints on hardware, software and/or programming language adopted by it.

The SCATTER PHY contains the following set of main features:

- Bursty transmissions: with discontinuous transmissions, it is possible to improve the use of an available spectrum band and to coordinate its usage with other networks/radios in an opportunistic/intelligent/collaborative way.

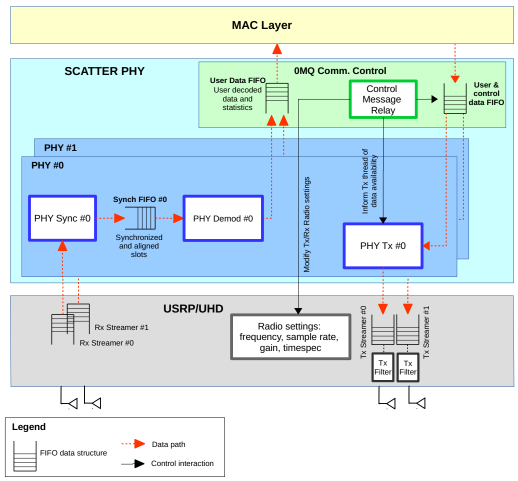

Figure 1: High-level architecture of the SCATTER PHY.

- Dual-Concurrent PHYs: having two physical interfaces simultaneously transmitting and receiving at independent frequencies enables Multi-Concurrent-Frequency Time-Division Multiple Access (McF-TDMA) scheme to be implemented by the MAC layer. This ability for concurrent allocations allows for a smarter spectrum utilization as vacant disjoint frequency chunks can be concurrently used.
- FPGA-based filtered transmissions: filtering the transmitted signal effectively minimizes out-of-band emissions (OOBE), which allows for better spectrum utilization, by enabling radios to have their transmissions closer to each other in the frequency domain, reducing spectrum wastage.
- Out-of-Band Full-Duplex operation: both PHYs operate completely independently, meaning that Tx and Rx modules are able to transmit and receive at different channels, set different gains and use different PHY BWs.
- Timed-commands: this feature allows the configuration of the exact time in the future to (i) start a transmission and (ii) change Tx/Rx frequencies/gains. This allows the MAC layer to implement a TDMA scheme.

Figure 1 illustrates the different software/hardware layers composing the SCATTER PHY and the threads within each one of them. Red dashed arrows indicate data paths while black arrows indicate control/information interaction between threads.

SCATTER PHY is a discontinuous transmission-based PHY, which transmits data bursts in small transport units called subframes. A subframe is a container through which user and control data is exchanged in the network. The SCATTER PHY is built upon the srsLTE library [6], and therefore absorbs and evolves on top of the existing LTE features. srsLTE is a free and open-source LTE software library developed by Software Radio Systems (SRS) [6]. We adopt Orthogonal Frequency-Division Multiplexing (OFDM) as the SCATTER PHY waveform. OFDM is a mature technology, widely implemented in a variety of products due to its several advantages, such as robustness to severe multipath fading, low implementation complexity, easy integration with MIMO, facilitated channel estimation, etc. [10]. User and control data are mapped into subcarriers over 14 OFDM symbols spanning 1 ms. The control data signal carries the used modulation and coding scheme (MCS) for the current transmission and the number of subsequent subframes with user data modulated with that specific MCS.

The control data signal is used at each one of the PHY receivers to automatically detect the number of allocated RBs, the location of the allocated RBs in the resource grid, and the MCS used to transmit data of a specific user. By embedding the control information into the transmitted signal, the Medium Access Control (MAC) layer does not need to know in advance the number of subframes and MCS in a given channel occupancy time (COT). Upon correct user data decoding, each PHY informs the MAC layer of the number of received bytes, the corresponding MCS, and decoding statistics (e.g. RSSI, SINR, number of detection/decoding errors, etc.).

SCATTER PHY employs discontinuous (i.e., bursty) transmission of subframes. A subframe is the basic transmission unit of the SCATTER PHY and each one is 1 ms long. SCATTER PHY works with two types of subframes, namely, synchronization and data-only subframes. A synchronization subframe carries the synchronization signal, reference signal, control, and user data. A data-only subframe carries the reference signal and user data. The synchronization signal is a 72 symbols long sequence that is generated using Zadof-Chu sequences [11]. The control signal is based on maximum length sequences (M-sequences) [13], where two M-sequences, each of length 31, are used to transmit the information necessary to decode the user data. The reference signal is used to estimate the channel and then equalize the received signal so that channel interference to the desired signal is minimized.

The SCATTER PHY allows bursty transmissions with variable COT, i.e., the number of subframes to be transmitted in sequence without any gap (i.e., idle time) between them is variable. The number of subframes in a COT is derived based on MCS, PHY BW and data length (i.e., number of bits to be transmitted) parameters sent by the MAC layer in the control message. The minimum COT is equal to 1 ms and is equivalent to the synchronization subframe. Variable COT enables the support of different traffic loads and channel occupancy. Every subframe can carry a pre-defined number of bits, which is based on the MCS and current PHY BW.

Each one of the PHY modules is split into three submodules, namely, PHY Tx, PHY Rx Synchronization and PHY Rx Demodulation where each one of them runs on an exclusive, standalone thread. The reason for having a multi-threaded PHY implementation is that it allows independent critical and/or time-consuming tasks to be executed simultaneously (i.e., concurrently), which improves computing performance and efficiency. Allied with multi-core enabled Central Processing Units (CPUs), the multi-threaded PHYs naturally support full-duplex communications mode, i.e., each one of the two PHYs can simultaneously transmit and receive at different frequencies, which consequently results in higher throughput. The PHY Tx threads (#0 and #1) are responsible for modulation and transmission of data (i.e., user and control data). The PHY Rx Synchronization threads are responsible for the detection of the synchronization (Synch) signal, decoding of the control data, carrier frequency offset (CFO) estimation/correction and subframe time-alignment tasks. Detection of the Synch signal is carried out through a two-stage detection algorithm, which at the first stage correlates the received signal with a locally-stored version of the synchronization subframe with no data and control signals. If the peak-to-side-lobe ratio (PSR) is greater than a constant threshold, then, the second stage applies a Cell-Average Constant False Alarm Rate (CA-CFAR) algorithm to the OFDM symbol carrying the Synch signal [11]. The two-stage approach employed by the SCATTER PHY improves the Synch signal detection when compared to a detection approach that only uses the PSR of the correlation calculated at the first stage.

The CFO estimation task is split into coarse and fine estimations/corrections, where the coarse estimation is based on the Synch signal and the fine estimation is based on the Cyclic Prefix (CP) portion of the OFDM symbols [5]. The integer part of the frequency offset (i.e., integer multiples of the subcarrier spacing) is estimated and corrected by the coarse CFO algorithm, which is based on the maximization of the correlation of the received synchronization signal with several locally generated frequency offset versions of it. On the other hand, the fractional frequency offset (i.e., offset values less than one half of the subcarrier spacing) is estimated and corrected by the fine CFO algorithm, which is based on the phase difference of the correlation between the CP and the last part of the OFDM symbol (i.e., the portion used to create the CP).

The PHY Rx Demodulation threads take care of user data demodulation, i.e., OFDM demodulation (FFT processing and CP removal), channel estimation/equalization, resource demapping, symbol demodulation, de-scrambling, de-interleaving/de-rate matching, turbo decoding, de-segmentation and, CRC checking. Each PHY receives data and control messages from the ZeroMQ Data/Control module. Decoded user data and statics related to the PHY operation (Rx/Tx statistics) are sent directly to the upper layers through the 0MQ bus. Regarding numerology, the SCATTER PHY uses subcarrier spacing of 15 kHz, a CP of 5.2 μs and supports 1.26, 2.7, 4.5, and 9 MHz bandwidths. Each PHY channel bandwidth (BW) can be set through the command line at start up, and changed in real-time through Tx control messages.

SCATTER PHY frame structure is depicted in Figure 2. As can be seen, synchronization and control data signals are added only to the very first subframe of a COT. With this frame structure, synchronization (i.e., detection of synchronization signal, time-alignment and CFO estimation/correction) and control data decoding only happen for subframe #0, i.e., the synchronization subframe.

The Payload Data Unit (PDU) adopted by SCATTER PHY is a Transport Block (TB). Therefore, a TB is the payload coming from the MAC layer and given to the individual PHYs to be encoded and transmitted over the air. One TB consists of a number of bits that can be accommodated within a 1 ms long subframe given the selected PHY BW and MCS. Therefore, given the PHY BW and the desired MCS, the MAC layer can find the number of bits that can be handled by a 1 ms long subframe. Table I presents the coding rate for each one of the 32 defined MCS values.

TABLE I: MCS VALUES AND THEIR RESPECTIVE CODE RATE.

		Code Rate			
MCS	**Modulation**	**1.26 MHz**	**2.7 MHz**	**4.5 MHz**	**9 MHz**
0	QPSK	0.0857	0.0900	0.0940	0.0940
1	QPSK	0.1143	0.1200	0.1220	0.1220
2	QPSK	0.1357	0.1467	0.1480	0.1480
3	QPSK	0.1714	0.1933	0.1920	0.1920
4	QPSK	0.2107	0.2400	0.2400	0.2440
5	QPSK	0.2571	0.2933	0.2960	0.2920
6	QPSK	0.3071	0.3400	0.3440	0.3440
7	QPSK	0.3571	0.4000	0.4160	0.4130
8	QPSK	0.4071	0.4667	0.4640	0.4690
9	QPSK	0.4714	0.5200	0.5360	0.5330
10	16QAM	0.2357	0.2600	0.2680	0.2665
11	16QAM	0.2571	0.2933	0.2920	0.2905
12	16QAM	0.3000	0.3267	0.3280	0.3305
13	16QAM	0.3357	0.3667	0.3800	0.3825
14	16QAM	0.3857	0.4267	0.4290	0.4298
15	16QAM	0.4286	0.4667	0.4770	0.4718
16	16QAM	0.4429	0.5000	0.5090	0.5078
17	64QAM	0.2952	0.3333	0.3393	0.3385
18	64QAM	0.3238	0.3600	0.3553	0.3625
19	64QAM	0.3524	0.3911	0.4033	0.4087
20	64QAM	0.3905	0.4411	0.4353	0.4407
21	64QAM	0.4286	0.4678	0.4727	0.4727
22	64QAM	0.4571	0.5122	0.5047	0.5100
23	64QAM	0.4952	0.5478	0.5570	0.5642
24	64QAM	0.5333	0.5833	0.5970	0.6042
25	64QAM	0.5714	0.6189	0.6210	0.6308
26	64QAM	0.5905	0.6633	0.6690	0.6770
27	64QAM	0.6190	0.6900	0.7010	0.7010
28	64QAM	0.6571	0.8056	0.8067	0.8178
29	64QAM	0.6952	0.8322	0.8280	0.8552
30	64QAM	0.7333	0.8617	0.8493	0.8925
31	64QAM	0.7714	0.8883	0.8707	0.9240

The communication between the SCATTER PHY and the MAC layer is carried out through the exchange of four predefined messages. The first two, namely, Tx and Rx Control messages, are used to manage subframe transmission and reception respectively. The parameters carried by these two messages can be configured and sent to the individual PHYs by the MAC layer before the transmission of every subframe, hence allowing runtime configuration. The other two messages, namely, Tx and Rx statistics messages, are used to provide real-time feedback from each PHY to the MCA layer, yielding vital information necessary for such layer to take actions.

Tx control messages carry the user data (i.e., TB) to be transmitted and Tx parameters related to that transmission, namely, PHY ID, MCS, data length, Tx gain, Tx channel, Tx PHY BW, and transmission timestamp. The transmission timestamp parameter enables time-scheduled transmissions,

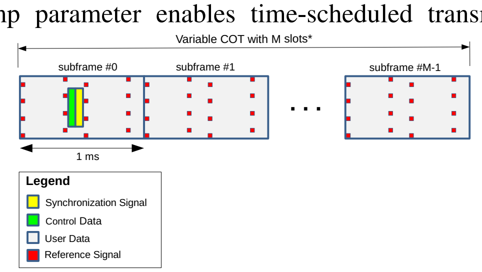

Figure 2: SCATTER PHY frame structure.

(a) (b)

Figure 3: Comparison of McF-TDMA feature. (a) FPGA-based filters disabled. (b) 128 order FPGA-based FIR filter enabled.

which allows the MAC layer to implement a Multi-Frequency (MF) Time Division Multiple Access (TDMA) medium access scheme. **Rx control** messages are used to configure Rx channel, Rx gain, and Rx PHY BW of a specific PHY, which addressed through the PHY ID parameter. The PHY ID parameter is used to specify to each one of the two PHYs a control message is meant to.

The other two messages, namely, Rx and Tx statistics, are used to inform the MAC layer of the Rx and Tx processing results of each PHY, respectively. **Rx statistics** messages carry the PHY ID, received data and reception statistics related to the received data such as Channel Quality Indicator (CQI), Received Signal Strength Indication (RSSI), decoded MCS, subframe error counter, decoding time, etc. **Tx statistics** messages inform upper layers of transmission statistics like coding time, the total number of transmitted subframes of each specific PHY. Table II summarizes all the real-time configurable parameters and statistics offered by the SCATTER PHY.

[1]Depends on the USRP daughter board installed [9].
[2]Bandwidths: 1.26, 2.7, 4.5, and 9 MHz respectively.

TABLE II: REAL-TIME CONFIGURABLE PARAMETERS AND STATISTICS.

Message	Parameter	Type	Unit	Range
Tx control	PHY ID	uint32	-	0-1
	MCS	uint8	-	0-28
	Tx gain	uint32	dB	depends on HW [1]
	Tx channel	uint32	-	≥ 0
	Tx PHY BW	uint8	MHz	0-3 [2]
	Transmission timestamp	uint64	s	≥ 0
	User Data length	uint32	-	> 0
	User data	uchar[]	-	uchar range
Rx control	PHY ID	uint32	-	0-1
	Rx channel	uint32	-	≥ 0
	Rx gain	uint32	dB	depends on HW [1]
	Rx PHY BW	uint8	MHz	0-3 [2]
Rx statistics	PHY ID	uint32	-	0-1
	CQI	uint8	-	0-15
	RSSI	float	dBW	float range
	Noise	float	dBW	float range
	Decoded MCS	uint8	-	0-28
	Subframe error counter	uint32	-	≥ 0
	Decoding time	uint32	ms	≥ 0
	Data length	uint32	-	≥ 0
	Received data	uchar[]	-	uchar range
Tx statistics	PHY ID	uint32	-	0-1
	Coding time	uint32	ms	≥ 0
	Number of transmitted subframes	uint32	-	≥ 0

A. COT-based Filtering

OFDM-based waveforms are not suited for spectral coexistence due to their poor spectral localization [7]. This problem is caused by the rectangular pulse-shape used in OFDM, which leads to a sync-pulse property in the frequency domain with a very low second lobe attenuation of -13 dB [8]. Therefore, in order to guarantee a better spectral localization, i.e., lower out-of-band (OOB) emissions, and maintain the complex-domain orthogonality of the OFDM symbols is to apply some sort of filtering to the time domain subframes. The filtering process is applied to each COT of each PHY independently. The subframes comprising a COT are generated at the SW level and then filtered at HW level, by a 128 order FPGA-based FIR filter. The COT-based filtering improves closer coexistence with other radios (either belonging to our team or others), allowing radio transmissions to be closer in frequency. The filter used in the SCATTER PHY is designed and explained in [12].

This co-design SW/HW is used so that fast-processing high order filters can be implemented adding up very low latency to the transmission chain and still allowing the flexibility of the software-defined PHYs. The filter's coefficients applied to the COT are automatically selected according to the configured Tx PHY BW (i.e., the coefficients are selected in real-time based on the Tx PHY BW field in the Tx control message) as it needs to have its cut-off frequency changed to exactly filter the desired signal's bandwidth.

B. Benefits of the SCATTER PHY

The first advantage of the proposed PHY is that the COT-based filtering makes the SCATTER PHY more spectral efficient as the OOB emissions are reduced. OOB emissions interfere with nearby radios (i.e., radios with nearby channels), decreasing the quality of their received signals, which consequently impacts the throughput experienced by those radios. The reduced OOB emissions make SCATTER PHY ideal for coexistence with other radios, allowing it to operate closer to them in the frequency domain and consequently, reducing spectrum wastage while increasing the spectral efficiency. The second advantage offered by SCATTER PHY is the possibility to configure in real-time all PHY parameters through the control messages.

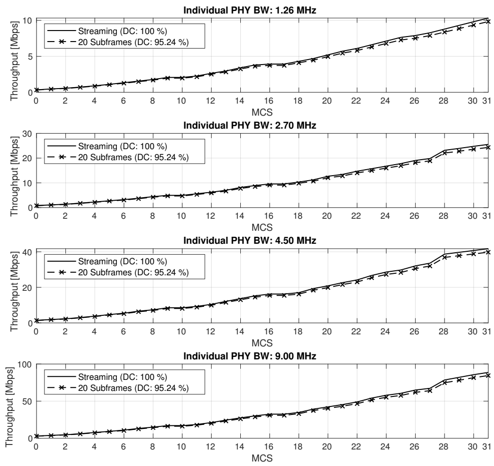

Figure 4: Throughput for different bandwidth, MCS, and DC.

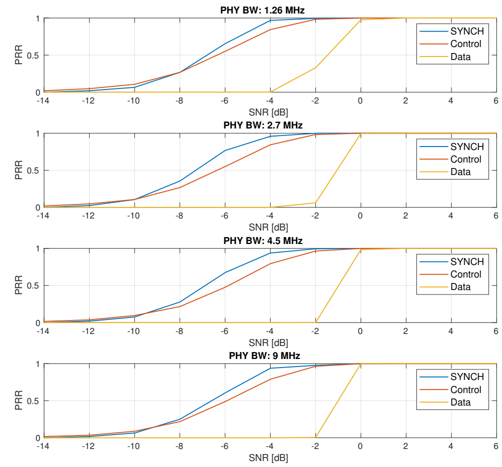

Figure 6: PRR comparison for the different signals.

III. Experiment Results

In this section, we present some experimental results in order to demonstrate the effectiveness and usability of the SCATTER PHY. All the experiments presented here were carried out with the framework running on servers with Intel Xeon E5-2650 v4 CPUs (@2.2 GHz,30 M cache, 9.60 GT/s QPI, Turbo, HT, 12 Cores/24 Threads, 105 Watts) with 128 GB of RAM memory connected to x310 USRPs with 10 Gigabit Ethernet links, and equipped with CBX-120 RF daughterboards [14].

Figure 3 shows two spectrograms collected during 40 ms over a 31.25 MHz bandwidth with both PHYs set to operate concurrently, at each instant, at two of six 4.5 MHz channels. This experiment intends to show that the SCATTER PHY is able to generate two concurrent and independent transmit channels per node. In this experiment, each PHY transmits a random number of subframes at randomly selected channels. The channel number and the number of transmitted subframes, i.e., COT, are randomly selected between the ranges 0-5 and 1-3, respectively. Here, a gap of 1 ms between consecutive transmission is used. As can be noticed in both figures, (a) and (b), SCATTER PHY is able to independently transmit at two distinct channels with a different number of subframes. Figure 3 (a) shows the case where no filtering is enabled. As can be noticed, OOBE might cause interference to adjacent channels and consequently decrease other radios' throughput. Figure 3 (b) shows the case when the FPGA-based FIR filters are enabled. As can be seen, when filtering is enabled, the OOBE is mitigated and consequently, interference to and from adjacent channels is also mitigated. Additionally, another consequence of the filtering use is that the channel spacing could be made smaller and consequently decrease the spectrum wastage.

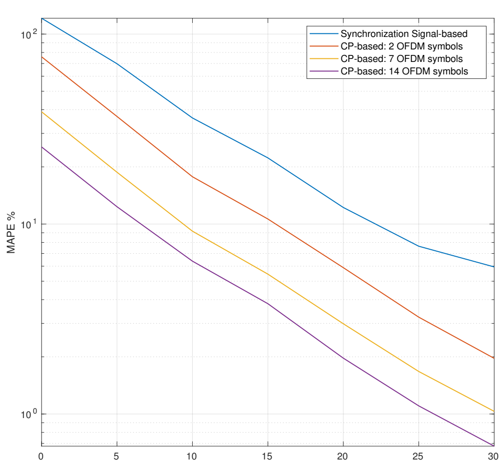

Figure 5: Comparison of CFO estimation error.

Figure 4 depicts throughput measurements taken with the SCATTER PHY for several PHY BW, MCS, and Duty Cycle (DC) values working in full-duplex mode (i.e., the two independent PHYs are simultaneously transmitting and receiving). We employ the full-duplex mode in order to check if it impacts somehow the measured throughout once in full-duplex mode the SCATTER PHY is being fully utilized. The measurement is taken for a COT of 20 ms with a gap of 1 ms between subsequent transmissions. This means a DC of 95.24 %. The throughput is averaged over 10 measurement intervals of 10 seconds each. During one measurement interval, the number of received bits is counted and then divided by the interval to calculate the throughput measured during that interval. The SNR on the link was set to 30 dB so that the packet reception rate for all MCS values was equal to 1. In the figure, for comparison reasons, we also add the theoretical maximum throughput achieved by the *Streaming* mode, where there is no gap between subsequent transmissions, i.e., a DC equal to 100 %. The theoretical maximum throughput is calculated dividing the TB size in bits for each MCS by 1 ms. As expected, the measured SCATTER PHY's throughput is close to the theoretical maximum throughput for all MCS and PHY BW values, yielding more than 84 Mbps for a BW of 9 MHz and MCS 31. Additionally, the operation in full-duplex mode has no visible impact on the achieved throughput. This is due to the powerful server, with 12 cores, used to run the SCATTER PHY.

Figure 5 compares the mean absolute percentage error (MAPE) of the fine CP-based CFO estimation method against estimation based on the synchronization signal. The results were obtained by connecting the Tx port to the Rx port of the same USRP so that frequency offset caused by the HW is minimal or nonexistent and adding the desired frequency offset plus AWGN at the SW level just before the subframes are transmitted. As can be seen, the fine CFO estimation algorithm, which is based on the CP portion of the OFDM symbols, outperforms the synch-based CFO algorithm even when only two consecutive CPs are averaged. Additionally, we see that the performance of the CP-based estimation improves as the number of averaged CPs increases, however, the downside of averaging more CPs is an increase in the processing time. The CP-based CFO estimation method employed by the SCATTER PHY is an improvement over the CFO estimation implemented by the srsLTE. The SCATTER PHY implementation makes use of single-instruction multiple data (SIMD) instructions in order to decrease the CFO estimation and correction computational time allowing it to use more number of averaged CPs. Additionally, SCATTER PHY implements a more accurate and fine-grained complex exponential signal for generating the correction signal, making it possible to generate signals with frequencies that are closer to the estimated ones.

Figure 6 presents the comparison of the packet reception rate (PRR) for each one of the signals carried by a SCATTER PHY subframe, namely, synchronization, control and data signals over several SNR and PHY BW values. The SNR is calculated based on the power of a 1 ms subframe, therefore, before adding noise to the transmitted subframe, the subframe power is calculated and then the necessary noisy power to achieve the desired SNR is calculated. The MCS used for modulating the user data is set to 0, which is the most robust coding scheme, allowing SCATTER PHY to decode data in low SNR scenarios. The purpose of this experiment is to identify the lowest possible SNR at which SCATTER PHY can still correctly decode the user data. This experiment is run by adding a channel emulator between the Tx and Rx sides of a single PHY instance. At the Tx side, the generated subframes, instead of being sent to the USRP HW are sent to an abstraction layer that emulates the HW and adds AWGN noise to the transmitted signal, next abstraction layer transfers the noisy signal to the receiving side of the PHY. The PRR is averaged over 10^4 trials, where at each trial, the Tx side of the PHY sends a single synchronization subframe. As can be seen, synchronization and control signals have a better PRR performance than that of the data decoding, however, SCATTER PHY's PRR performance is limited by the ability of correctly decoding the data section of a subframe. Additionally, we see that the Data PRR is better for the 1.26 MHz case, which is due to the fact that compared to the other PHY BW values, MCS 0 for the 1.26 MHz case, carries more redundancy bits, as shown by Table I, making it more robust against noise. The other two signals, synchronization, and control present similar PRR curves for all PHY BW values. Moreover, it is noticeable that the data PRR is equal to 1 for SNR values greater than or equal to 0 dB.

IV. Conclusion

The paper has described the design, features, and benefits of the SCATTER PHY. The main features include efficient interaction with higher layers of the system, ability to transmit/receive two different types of packets which contain synchronization signals, control data, and user data. The PHY layer communication operates with various modulation and coding schemes, several different channel bandwidths, and automatically improves the CFO correction by utilizing the OFDM CP in addition to the synchronization subframe detected in the signal. SCATTER PHY minimizes OOBE by utilizing custom-built FPGA-based FIR filters on the transmitter side, configured for each channel bandwidth at run time. Additionally, two independent physical interfaces are available, which effectively enable usage of full-duplex and theoretically double the possible throughput achieved with one interface. The described benefits were verified experimentally. The verification shows that using full-duplex makes no impact on the achieved throughput in each interface. Spectrometer measurement results showcase the capabilities of the FPGA-based FIR filter, resulting in negligible OOBE, which in turn increases channel utilization and facilitates communication in neighboring channels by other radios. Using CP-based CFO estimation and correction is demonstrated to outperform the sync-signal based correction in all cases, at the expense of higher processing time. Finally, the packet reception rate for all types of transmitted signals is presented, showing that the SCATTER PHY can successfully decode user data at SNR values as low as 0 dB.

Acknowledgment

This work was funded by the European Union's Horizon 2020 research and innovation programme under grant agreement No. 732174 (ORCA project).

References

[1] Google, *protobuf - Protocol Buffers - Google's data interchange format*, [online]. Available: http://code.google.com/p/protobuf/, 2011.

[2] P. Hintjens, *ZeroMQ Messaging for Many Applications*, O'Reilly Media, March 2013.

[3] Ettus Research LLC, *USRP X Series*, [online]. Available: https://www.ettus.com/product/category/USRP-X-Series

[4] National Instruments, *Overview of the NI USRP RIO Software-Defined Radio*, White Paper, December 2015. [online]. Available: http://www.ni.com/white-paper/52119/en/

[5] Feng Wang, and Yu Zhu, *An efficient CFO estimation algorithm for the downlink of 3GPP-LTE*, International Conference on Wireless Communications and Signal Processing (WCSP), November 2011.

[6] I. Gomez-Miguelez, et. al., *srsLTE: An Open-Source Platform for LTE Evolution and Experimentation*, ACM WiNTECH Workshop, October 2016.

[7] Quentin Bodinier, Faouzi Bader, and Jacques Palicot, *On Spectral Coexistence of CP-OFDM and FB-MC Waveforms in 5G Networks*, IEEE Access, vol. 5, July 2017.

[8] Hao Lin, and Pierre Siohan, *Major 5G Waveform Candidates: Overview and Comparison*, Book chapter in Signal Processing for 5G, Wiley, August 2016.

[9] Ettus Research LLC, *Products*, [online]. Available: https://www.ettus.com/product

[10] Sravanthi Kanchi, et. al., *Overview of LTE-A technology*, IEEE Global High Tech Congress on Electronics (GHTCE), March 2014.

[11] Felipe A. P. de Figueiredo, et. al., *Multi-stage Based Cross-Correlation Peak Detection for LTE Random Access Preambles*, Revista Telecomunicacoes, vol. 15, September 2013.

[12] Felipe A. P. de Figueiredo et. al., *A Spectrum Sharing Framework for Intelligent Next Generation Wireless Networks*, IEEE Access, vol. 6, pp. 60704-60735, December 2018.

[13] D. V. Sarwate, and M. B. Pursley, *Crosscorrelation properties of pseudorandom and related sequences*, Proc. IEEE, vol. 68, no. 5, pp. 593-619, May 1980.

[14] Ettus Research LLC, *CBX 1200-6000 MHz Rx/Tx (120 MHz, X Series only)*, [online]. Available: https://www.ettus.com/product/details/CBX120

[15] Ettus Research LLC, *USRP Hardware Driver*, [online]. Available: http://files.ettus.com/manual/