A TDD Distributed MIMO Testbed Using a 1-Bit Radio-Over-Fiber Fronthaul Architecture

(article starts on next page)

A TDD Distributed MIMO Testbed Using a 1-Bit Radio-Over-Fiber Fronthaul Architecture

Lise Aabel, Sven Jacobsson, Mikael Coldrey, Frida Olofsson,
Giuseppe Durisi, *Senior Member, IEEE,* and Christian Fager, *Senior Member, IEEE*

Abstract—We present the uplink and downlink of a time-division duplex distributed multiple-input multiple-output (D-MIMO) testbed, based on a 1-bit radio-over-fiber architecture, which is low-cost and scalable. The proposed architecture involves a central unit (CU) that is equipped with 1-bit digital-to-analog and analog-to-digital converters, operating at 10 GS/s. The CU is connected to multiple single-antenna remote radio heads (RRHs) via optical fibers, over which a binary RF waveform is transmitted. In the uplink, a binary RF waveform is generated at the RRHs by a comparator, whose inputs are the received RF signal and a suitably designed dither signal. In the downlink, a binary RF waveform is generated at the CU via bandpass sigma-delta modulation. Our measurement results show that low error-vector magnitude (EVM) can be achieved in both the uplink and the downlink, despite 1-bit sampling at the CU. Specifically, for point-to-point over-cable transmission between a single user equipment (UE) and a CU equipped with a single RRH, we report, for a 10 MBd signal using single-carrier 16QAM modulation, an EVM of 3.3% in the downlink, and of 4.5% in the uplink. We then consider a CU connected to 3 RRHs serving over the air 2 UEs, and show that, after over-the-air reciprocity calibration, a downlink zero-forcing precoder designed on the basis of uplink channel estimates at the CU, achieves an EVM of 6.4% and 10.9% at UE 1 and UE 2, respectively. Finally, we investigate the ability of the proposed architecture to support orthogonal frequency-division multiplexing (OFDM) waveforms, and its robustness against both in-band and out-of-band interference.

Index Terms—Distributed MIMO, 1-bit sampling, radio-over-fiber.

I. Introduction

THE demands for more uniform quality of service, lower latency, higher reliability, and improved energy efficiency are driving the wireless network design towards 6G [5]. Distributed multiple-input multiple-output (D-MIMO) is a promising technology to satisfy these demands, see, e.g., [6]. In D-MIMO, a central unit (CU) is connected to multiple spatially separated antenna units, which we shall refer to as remote radio heads (RRHs). The advantages of D-MIMO over traditional co-located MIMO include enhanced robustness towards shadow fading, reduced spatial correlation between channels, and increased energy efficiency [7]. From a network perspective, D-MIMO is a key technology to implement the cell-free massive MIMO network architecture described in, e.g., [8]–[10].

The theoretical principles underpinning the optimal design of D-MIMO architectures are well understood [10]. However, many issues remain concerning the practical realization of this technology. One challenge in the implementation of D-MIMO is to ensure phase-synchronized transmission and reception between the RRHs, which is necessary to serve simultaneously multiple user equipment (UEs) via spatial processing. In theoretical investigations, it is often assumed that frequency up- and down-conversion are performed locally at the RRHs, and that the fronthaul link connecting the RRHs to the CU carries samples of the baseband signals. One way to achieve this is to equip each RRH with a local oscillator. To guarantee phase-coherent transmission and reception, the local oscillators must then be synchronized. This is challenging in the hardware domain, and costly in the digital domain, in terms of signaling overhead.

To bypass this issue, we consider in this paper an architecture in which up- and down-conversion are performed digitally at the CU. This eliminates the issue of synchronizing local oscillators. The drawbacks are that radio-frequency (RF) signals, rather than baseband signals, need to be exchanged over the fronthaul link and that analog-to-digital (ADC) and digital-to-analog (DAC) converters need to operate at a much higher sampling rate. To alleviate both problems, we focus on an architecture in which the fronthaul links consist of optical fibers, and the RF signals are mapped into a two-level (binary) optical waveform, prior to transmission over the fiber. In the following, we shall refer to this architecture as *1-bit radio-over-fiber fronthaul*. This architecture is advantageous in terms of cost, power consumption, and scalability [11]. Indeed, the use of binary optical waveforms allows us to leverage the recent development of low-cost, high-speed optical components, driven by the needs of, e.g., data centers. Also, it allows us to use power-efficient 1-bit ADCs and DACs at the CU.

An alternative architecture based on analog radio-over-fiber fronthaul has also been considered in the literature for D-MIMO applications, see, e.g., [12], although no phase-coherent downlink transmission based on uplink channel estimation has been demonstrated yet. In this alternative architecture, the RF signal is modulated directly onto the optical

TABLE I
UPLINK MEASUREMENTS USING 1-BIT RADIO-OVER-FIBER FRONTHAUL

Ref.	Type	AGC	Dither gen.	Channel	Bandwidth	Fronthaul rate
[1]	SISO	No	FPGA	Cable	10 MHz	10 Gbps
[2]	SISO	No	Sig. Gen.	Cable	5 MHz	NA
[3]	SISO	No	FPGA	Cable	20 MHz	10 Gbps
[4]	SISO	No	NA	Cable	5 MHz	3 Gbps
This work	Multi-user-D-MIMO	Yes	FPGA	Wireless	100 MHz	10 Gbps

carrier. This, however, makes this architecture less robust toward nonlinearities occurring in the optical domain [11].

A. Previous Work

D-MIMO architectures based on the 1-bit radio-over-fiber fronthaul have been studied previously in the literature [1]–[4], [13]–[17]. In the downlink, bandpass sigma-delta ($\Sigma\Delta$) modulation is used to map the RF signal into a two-level waveform [13]–[17]. At the CU, a 1-bit quantized and oversampled version of the RF signal is then transmitted to the RRHs via a two-level optical waveform. The desired RF signal is finally obtained at each RRH via a bandpass filtering operation. In the uplink, the RF signal is mapped into a two-level waveform at the RRHs by means of a comparator, which is fed with the RF signal and a suitably designed dither signal [1]–[4]. The resulting two-level waveform is transmitted to the CU via the fiber optical link, where it is sampled with 1-bit resolution.

Bandpass $\Sigma\Delta$-over-fiber for D-MIMO fronthaul was demonstrated in [13], [16] via point-to-point measurements. In [14], [15], [17], downlink transmission with multiple RRHs was demonstrated. However, since only a downlink testbed was available, the channel-state information needed for the design of the precoder was transferred from the UEs to the CU via a cable. In [1]–[4], point-to-point uplink reception was demonstrated. The demonstrations involved only over-cable measurements, in which an RF signal produced by a signal generator was connected directly to one of the two ports of the comparator. In [2], the dither signal was also produced by a signal generator, whereas in [1], [3] it was generated at the CU and conveyed to the RRH via the downlink fronthaul.

A theoretical analysis of the 1-bit quantized uplink is presented in [18], [19]. In [18], the authors show analytically and via simulations that the per-symbol error-vector magnitude (EVM) degrades in the high SNR regime when dithering is not used. The EVM degradation is caused by a high correlation between the signal and the quantization noise. This correlation can be mitigated via dithering, and satisfactory EVM performance at high SNR can be achieved. Specifically, both white Gaussian noise and uniformly distributed binary signals were considered in [18] as possible dither signals. In [19], the authors quantify the minimum fronthaul rate (sampling rate) required, for a D-MIMO architecture with 1-bit radio-over-fiber fronthaul and a fixed number of RRHs, to outperform in terms of EVM a conventional co-located massive MIMO architecture. Furthermore, the paper sheds light on the optimal dithering power for the case of Gaussian dithering signal.

B. Contributions

We present a time-division duplex (TDD) D-MIMO testbed with 1-bit radio-over-fiber fronthaul and multiple RRHs. Compared to previous work on 1-bit radio-over-fiber fronthaul for D-MIMO, our testbed implements both uplink and downlink, and supports downlink spatial multiplexing via reciprocity-based uplink channel estimation performed at the CU. Compared to [1]–[4], the RRH receiver is complemented with filters, a TDD switch, amplifiers, and automatic gain control (AGC). In particular, the AGC provides a dynamic range to the RRH receiver, which allows us to apply the same optimized dither signal to received signals having different power. We use point-to-point over-cable measurements to determine the dynamic range available in the uplink, the ability of the proposed architecture to support 5G waveforms, and also its robustness against both in-band and out-of-band interference.

We also present over-the-air measurements involving reciprocity calibration and downlink multi-user transmission based on uplink channel estimation, for a CU connected to 3 RRHs serving 2 UEs. Our measurements reveal that low EVM can be achieved in both uplink and downlink, despite the 1-bit quantization, using quantization-unaware linear spatial processing at the CU.

In Table I, we provide a summary of the uplink measurements conducted in [1]–[4] using the 1-bit radio-over-fiber fronthaul, as well as of the novel ones presented in this paper.

C. Paper Outline

The rest of this paper is organized as follows. In Section II, we detail the testbed architecture. In Section III, we discuss the testbed hardware and its operations. In Section IV, we present point-to-point over-cable measurements for downlink and uplink. In Section V, we discuss system reciprocity calibration. We present over-the-air multi-user D-MIMO measurements in Section VI. Finally, we provide some concluding remarks in Section VIII.

II. TESTBED ARCHITECTURE

A block diagram of the proposed D-MIMO testbed is provided in Fig. 1. In the figure, we consider the case in which a CU, connected to 3 RRHs, serves 2 UEs. The downlink and uplink are illustrated by the upper (marked in blue) and lower (marked in yellow) block chains, respectively. The signal generation and digital signal processing are performed at the CU, which is equipped with 1-bit DACs and ADCs, and is connected to the RRHs via optical fibers, carrying binary waveforms. In the next two subsections, we describe in detail the uplink and downlink blocks.
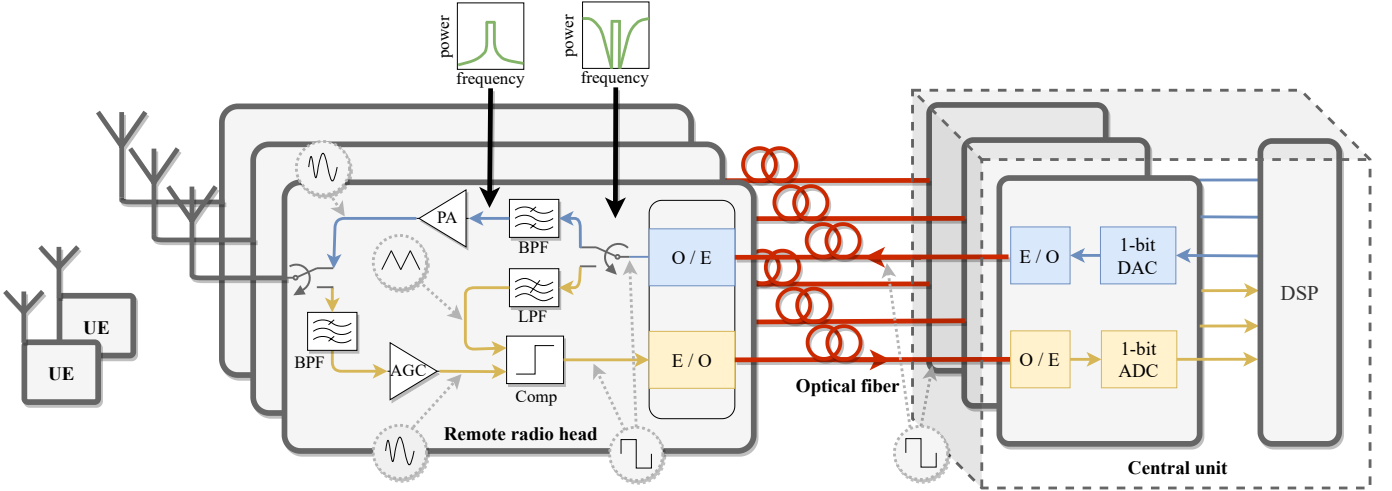
Fig. 1. Block diagram of the proposed TDD D-MIMO testbed. The downlink (upper block chain in blue) consists of the following blocks: digital signal processing (DSP), 1-bit DAC, electrical-to-optical converter (E/O), optical fiber, optical-to-electrical converter (O/E), BPF, and PA. The uplink (lower block chain in yellow) consists of the following blocks: BPF, AGC, comparator (Comp), E/O, optical fiber, O/E, 1-bit ADC, and DSP. The dither signal in the uplink is conveyed to the RRHs via the downlink path, and reconstructed at the RRH by means of a LPF. The frequency spectrum of the signals in different parts of the architecture is illustrated by the green plots, and the waveforms are indicated by the dashed circles.

A. Sigma-Delta-over-Fiber Downlink

At the CU, the coded information bits intended for each UE are first modulated and precoded. Then, after pulse shaping, the baseband signal intended to each RRH is digitally up-converted to RF and then transformed into an oversampled 1-bit stream by a bandpass $\Sigma\Delta$ modulator. This binary stream is converted to an analog signal by a 1-bit DAC, fed to an electrical-to-optical converter and sent to the corresponding RRH over a fiber optical cable. At each RRH, the received optical binary signal is converted to the electrical domain by means of an optical-to-electrical converter. Thanks to the properties of $\Sigma\Delta$ modulation, the RF signal is reconstructed at the RRHs simply by means of bandpass filtering. Finally, the resulting signal is passed through a power amplifier (PA) and fed to the antenna port.

The principle of $\Sigma\Delta$ modulation is to use oversampling and noise shaping to place the quantization noise introduced by the 1-bit DAC outside the bandwidth of the desired signal. As described in, e.g., [20], the amount of in-band quantization noise introduced by the $\Sigma\Delta$ modulator depends on two parameters: 1) the oversampling ratio

$$\text{OSR} = \frac{f_s}{2W}, \tag{1}$$

where f_s is the sampling rate at which the bandpass $\Sigma\Delta$ modulator operates and W is the bandwidth of the RF signal, and 2) the order of the $\Sigma\Delta$ modulator. Note that, the larger f_s, the lower the in-band quantization noise power. Furthermore, the order of the $\Sigma\Delta$ modulator, i.e., the number of integrators, determines the shape of the noise spectrum. The oversampling ratio and the order of the modulator can thus be adapted to achieve a desired in-band quantization noise power. However, higher-order modulators are more difficult to build, introduce more delay, consume more power, and are more prone to

instability. The parameters of the $\Sigma\Delta$ modulator implemented in the testbed are provided in Section IV.

B. Dithered 1-bit Quantization Uplink

In the uplink, the RF signal at the antenna port of each RRH is first passed through a bandpass filter (BPF), and then through an AGC. The AGC ensures that the amplitude of the filtered RF signal is within a desired range, to make dithering effective. A binary representation of the RF signal is then generated by feeding the output signal of the AGC and a dither signal to the comparator. As demonstrated in, e.g., [18], non-subtractive dithering can help whiten the quantization noise, which results in an improved EVM, especially in the high-SNR regime.

Similar to [2], [3], [21], [22], we use a baseband triangular waveform as a dither signal. In our architecture, the dither signal intended for each RRH is generated digitally at the CU. This allows us to adapt the amplitude and the frequency of the dither signal to match the characteristics of the received RF signal. This dither signal is then baseband $\Sigma\Delta$-modulated and conveyed to the RRHs via the downlink optical fibers, which are idle during the uplink in our TDD architecture. At the RRHs, the dither signal is reconstructed by means of a lowpass filter (LPF). The binary output from the comparator, which resembles a pulse-width modulated RF signal, is provided to the electrical-to-optical converter and sent to the CU, where it is fed to an optical-to-electrical converter and sampled by a 1-bit ADC. Digital signal processing in the uplink includes digital down-conversion, matched filtering, channel estimation, spatial combining, demodulation, and decoding.

III. HARDWARE

The D-MIMO testbed operates at a carrier frequency of 2.35 GHz, which is dictated by the choice of the analog BPFs.
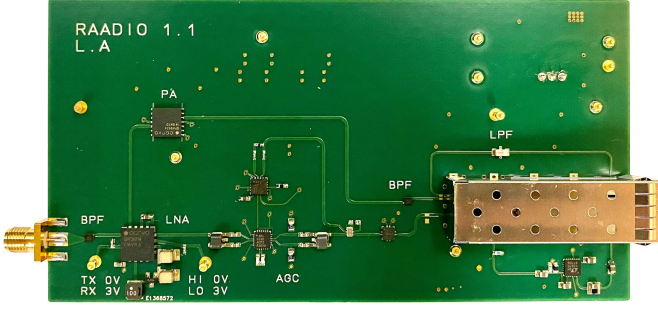
Fig. 2. An RRH printed circuit board. The following components/functions are marked on the board: BPF, PA, LNA, AGC, and LPF. The settings for the switch/LNA are marked as TX and RX and the gain modes are marked as HI and LO.

Specifically, the BPFs used in the testbed are bulk acoustic wave reconstruction filters with 100 MHz bandwidth, centered at 2.35 GHz. The CU consists of the field-programmable gate array (FPGA) evaluation board *Altera Stratix V GT Transceiver Signal Integrity Development Kit* [23] and a computer. The FPGA evaluation board operates at a sampling rate of 10 GS/s, and the digital ports are equipped with SMA connectors. The digital signal processing is performed offline on the computer using MATLAB. Specifically, $\Sigma\Delta$ modulation is implemented using the toolbox described in [20]. Electrical-to-optical and optical-to-electrical conversions are performed by the *Avago AFBR-709SMZ small form-factor pluggable+* (SFP+) optical transceivers [24]. Each SFP+ is equipped with an 850 nm vertical-cavity surface-emitting laser and photodetector, designed to support 10 Gigabit Ethernet. The optical fiber cables in the testbed are 30 m long, of type *Optical Multimode 4*. The electrical interface of the SFP+ offers differential input and output ports, which we use to implement the comparator. A picture of the assembled RRH is provided in Fig. 2. We next describe the components of the transmitter, depicted in Fig. 3, and the receiver, depicted in Fig. 4, in the RRH.

A. RRH Transmitter Components

The downlink two-level RF signal from the CU is received at the RRH at the non-inverting electrical output port of the SFP+. This signal is then passed through a reconstructing bulk acoustic wave BPF,[1] centered at 2.35 GHz and with 100 MHz bandwidth. The output of the filter is fed to a PA,[1] that has 35.8 dB linear gain within the 2.3–2.4 GHz frequency range. We use a switch connecting the antenna to either the transmitter or the receiver chain, to move between uplink and downlink operations. Specifically, we use a combined switch/low-noise amplifier (LNA) module. Switch synchronization is currently implemented in the testbed by applying a common DC power supply voltage to all RRHs. When the switch is in the transmit state (downlink mode), the RF signal is connected directly to the output port of the module. Before reaching the antenna port, the downlink signal is passed through another bulk acoustic wave BPF, which suppresses out-of-band emissions.

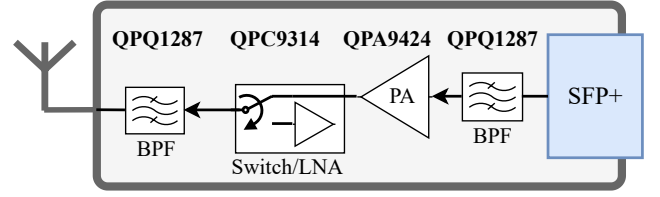[1]This component is the same as the one used in [14].

Fig. 3. Block diagram of the RRH transmitter components. It includes two identical BPFs, a switch/LNA module, a PA, and an SFP+ optical transceiver.

B. RRH Receiver Components

The RF signal received at the antenna port is filtered by the bulk acoustic wave BPF and connected to the switch/LNA module. When the switch is in the receive state (uplink mode), the RF signal is connected to the LNA within the module. This amplifier operates within the 2.3–2.7 GHz frequency range, provides a high gain mode of 33 dB or a low gain mode of 24 dB, and has a noise figure of 1.5 dB. We use only the low gain mode during the measurements presented in this paper. The AGC consists of a variable gain amplifier (VGA), which is controlled via a voltage applied to its gain control input port. The VGA operates within the 10 MHz–3 GHz range, provides a 45 dB dynamic range (from −30 dB to 15 dB), and has a noise figure of 12.5 dB. The voltage that controls the VGA is generated by a demodulating logarithmic amplifier, which acts as a power detector. The demodulating logarithmic amplifier converts the RF signal at its input to a voltage that is proportional to the RF signal power in dB. We use this demodulating logarithmic amplifier to maintain constant the average power of the signal at the output of the VGA at approximately −30 dBm. To do so, we use a 20 dB coupler to feed back a fraction of the RF signal at the output of the VGA to the demodulating logarithmic amplifier. The demodulating logarithmic amplifier is designed such that the adaptive controlling voltage does not change during 1 ms. After the AGC loop, i.e., at the output of the coupler, the RF signal is again amplified by an RF amplifier, to meet the input voltage requirements of the comparator.

The $\Sigma\Delta$-modulated triangular dither signal, sent from the CU via the fiber optical fronthaul, is received on the inverting output port of the SFP+ module and filtered by a 180 MHz LPF to remove out-of-band distortions. The resulting triangular waveform is then amplified by a 15 dB intermediate-frequency amplifier, operating within the 100 kHz–1 GHz range.

The output from the RF amplifier is fed to the SFP+ non-inverting input port and the output from the intermediate frequency amplifier is provided to the inverting input port. The SFP+ thus serves both as comparator and electrical-to-optical converter. Differential input voltages from 180 mV to 700 mV are supported by the SFP+.

IV. POINT-TO-POINT MEASUREMENTS

In this section, we present downlink and uplink over-cable point-to-point measurements. The purpose of these measurements is to investigate the RRH and fronthaul performance, without involving the wireless channel. In all measurements, we use 16QAM symbols, root-raised cosine pulse shaping with
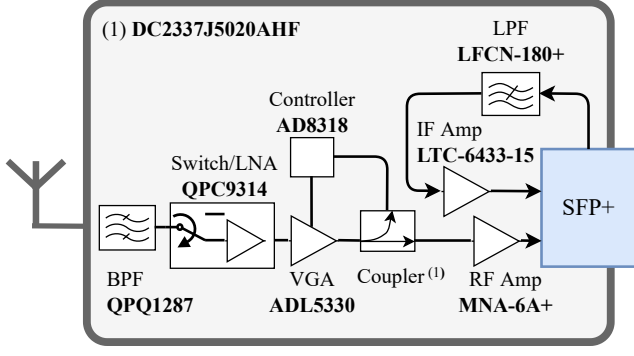
Fig. 4. Block diagram of the RRH receiver components. It includes a BPF, a switch/LNA module, a VGA, a coupler, a demodulating logarithmic amplifier, an RF amplifier, a LPF, an intermediate frequency (IF) amplifier, and an SFP+.
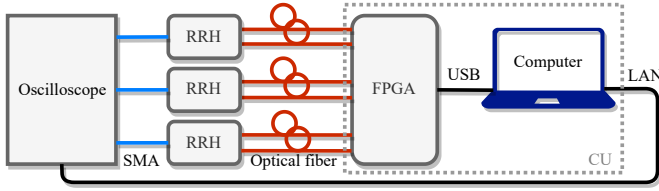

Fig. 5. Setup for downlink over-cable measurements. The CU consists of the computer and the FPGA evaluation board, which are connected via a universal serial bus (USB). The oscilloscope has a 4 GHz bandwidth and samples the signal at 10 GS/s.

a roll-off factor of 0.2, and 4^{th} order $\Sigma\Delta$-modulation. A $\Sigma\Delta$-modulator of order 4 strikes a good tradeoff between noise-spectrum shaping and implementation complexity [20].

A. Point-to-Point Downlink Over-Cable Measurements

We measure the downlink performance for a 10 MBd single-carrier RF signal in terms of EVM of the received constellation symbols, by using the setup described in Fig. 5. The antenna ports of the RRHs are connected to the input ports of the oscilloscope Rohde & Schwarz *RTO 1044*, which is connected to the computer via a LAN connection. We then compare the EVM measured for this setup, with the one measured for the case in which the FPGA is replaced by a pulse-pattern generator (PPG). A PPG provides high-quality binary waveforms with short rise time (8.5 ps). Our goal with these two measurements is to separate the noise generated by the RRHs and the fronthaul link, from the noise generated by the FPGA evaluation board.

First, a single-carrier RF signal with 10 MBd symbol rate is generated, bandpass $\Sigma\Delta$-modulated and written to the FPGA. The binary waveforms at the output ports on the FPGA evaluation board are transmitted over the optical links to the RRHs.

In Fig. 6(a) and 6(b), we present the baseband and passband power spectral densities of the signals received at the oscilloscope, when the same 10 MBd signal is passed through each of the three RRHs. We see from Fig. 6(a) that the power spectral densities are very similar. Furthermore, the low noise outside of the bandwidth of interest seen in Fig. 6(b) indicates that the noise suppression of the BPFs is adequate. We show
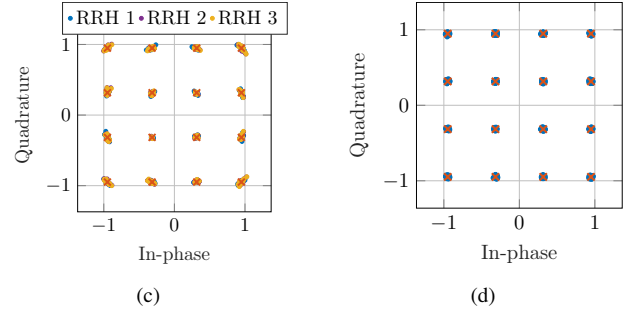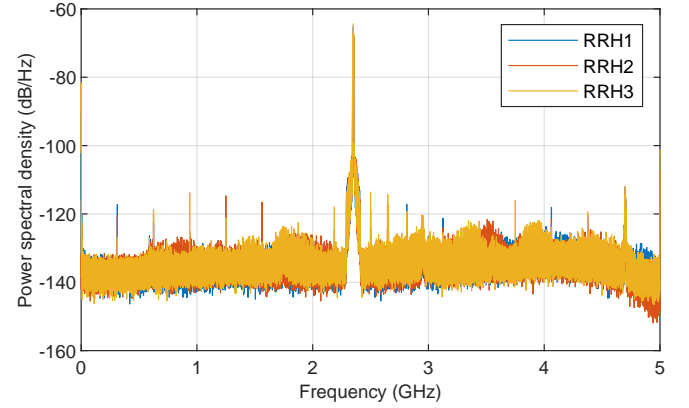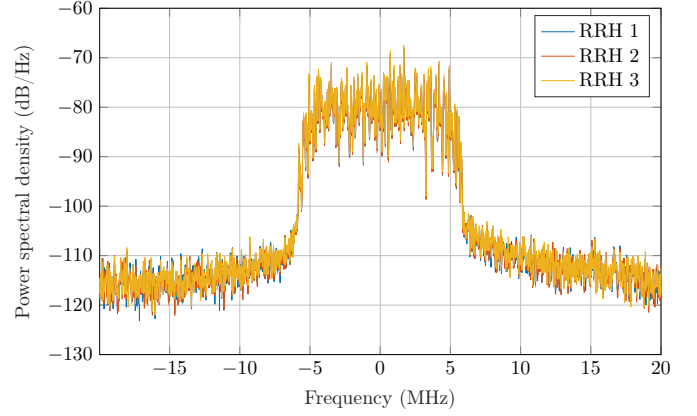

Fig. 6. (a) Welch power spectral density estimates of the in-band and adjacent bands when using the FPGA as binary waveform transmitter. (b) Welch power spectral density estimates from 0–$f_s/2$ of the RF signals when using the FPGA as binary waveform transmitter. (c) Constellation diagram obtained using 3 RRHs and FPGA. (d) Constellation diagram obtained using RRH 1 and PPG.

in Fig. 6(c) the corresponding received constellation diagrams, which results in an EVM of 3.3% for all RRHs.

In Fig. 6(d), we present the constellation diagram for the case in which the FPGA is replaced by the PPG and RRH 1 is used. The corresponding EVM is 2.3%. To understand the cause behind this EVM difference, we analyze next in detail the waveform generated by the FPGA and the PPG, corresponding to the bit pattern $[01010110]$ transmitted at 10 Gbps and repeated for 15 μs. The waveforms are recorded using a Keysight *UXR0334A* oscilloscope with a bandwidth of 33 GHz, which samples them at 128 GS/s. In Fig. 7 we provide
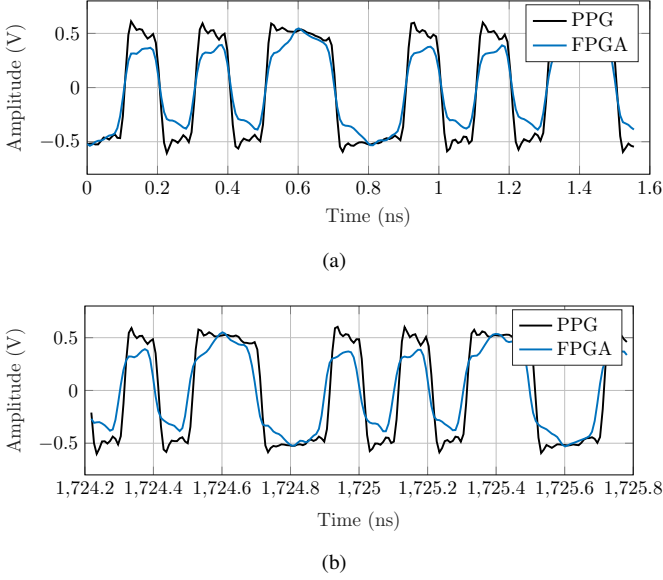
(a)

(b)

Fig. 7. Snapshot of FPGA and PPG output (a) at the first time instance and (b) at the second time instance.
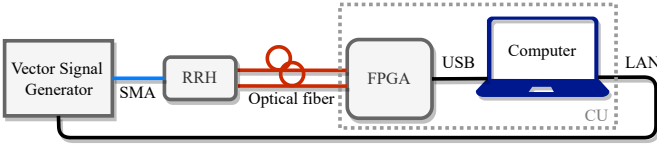

Fig. 8. Setup for uplink over-cable measurements.

a comparison between two snapshots of the two waveforms, taken at two different time instances. We see from Fig. 7(a), that, as expected, the pulses generated by the PPG have faster rise time. As a consequence, the short pulses generated by the FPGA have lower amplitude than the ones generated by the PPG. In Fig. 7(b), we present a second snapshot, recorded 1724 ns after the first snapshot. Compared to the first snapshot, we see that the FPGA signal is now also affected by an additional delay, which hints at a phase instability of the FPGA output. This phase instability as well as the amplitude reduction causes the additional noise observed in the constellation diagram in Fig. 6(c), compared to Fig. 6(d).

B. Point-to-Point Uplink Over-Cable Measurements

To measure the uplink performance in terms of the EVM of the received constellation symbols, we connect a Rohde & Schwarz *SMU 200A* vector signal generator via SMA cable to the antenna port of the RRH. The baseband signal is generated in MATLAB and provided to the vector signal generator, which performs frequency up-conversion. The measurement setup is shown in Fig. 8.

1) Dithering Signal: We investigate the effect on the EVM of changing the power and the frequency of the dither signal, when feeding the RRH with a 10 MBd signal. During this measurement, the average transmitted power is set to −40 dBm.

In Fig. 9, we present the EVM of the received constellation symbols, as a function of the frequency and the power of the dither signal. We observe that when the frequency of the
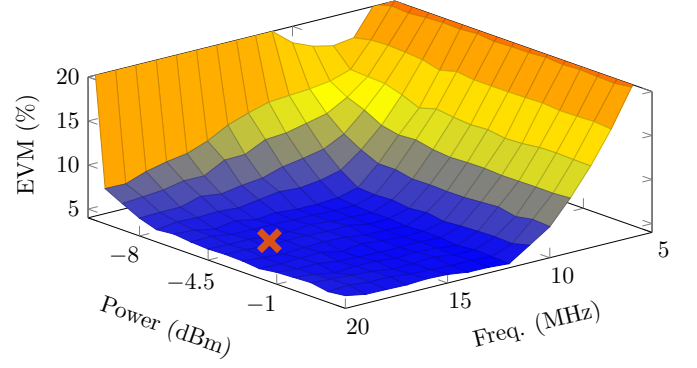

Fig. 9. EVM of the received symbol constellation as a function of the frequency and the power of the dither signal. The minimum EVM, marked in the plot with a cross, is 4.5%. It is achieved for a dither signal power of −4.5 dBm and frequency of 17 MHz.
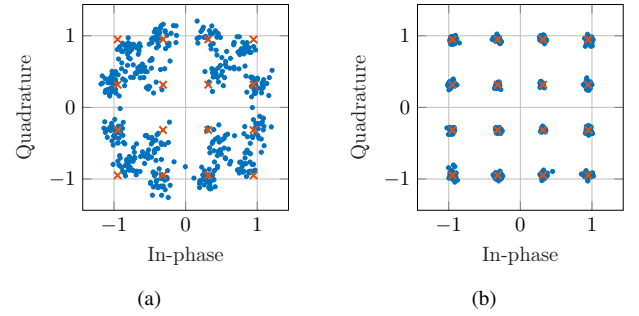

(a)   (b)

Fig. 10. The constellation diagram in (a) is obtained using a 5 MHz dither signal at −11.5 dBm and in (b) using a 17 MHz dither signal at −4.5 dBm.

dither signal is lower than the symbol rate of the transmit signal (i.e., 10 MBd) or when the power of the dither signal is below −8 dBm, the EVM is high. These results are in accordance with the ones reported in [3], [21]. Furthermore, we observe that changing the dither signal frequency and power has little effect on the EVM when the frequency is higher than 10 MHz and the power is higher than −8 dBm. The smallest measured EVM is 4.5% and occurs at 17 MHz and −4.5 dBm, but many combinations of dither signal frequency and power yield similar performance.

To illustrate the cause of the high EVM experienced when the choice of the power and frequency of the dither signal are suboptimal, we present in Fig. 10(a) the received constellation diagram for a poor choice of dither-signal power and frequency. We observe from the figure that the signal is severely distorted by the 1-bit quantization. In comparison, we illustrate in Fig. 10(b) the constellation diagram for a good choice of dither-signal power and frequency. We note that the EVM is greatly improved and that, differently from Fig. 10(a), the noise exhibits a circular symmetry around each constellation point.

2) Dynamic Range: We investigate the dynamic range of the RRHs by measuring the EVM of the received constellation symbols for different values of the transmit signal power and a fixed dither signal. The measurements are conducted using the same signal parameters as for the dither analysis.
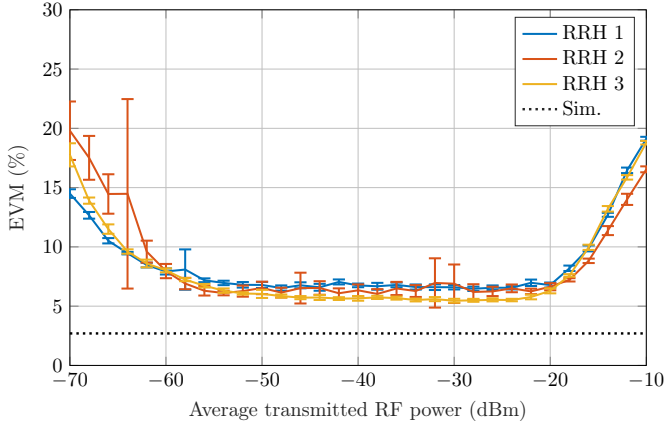
Fig. 11. Average per-symbol EVM for each RRH, versus the average transmitted RF power from the vector signal generator. For each point, we provide the standard deviation, based on ten measurements. The simulated EVM, obtained by assuming ideal hardware components, is marked with a dashed line.

Fig. 12. Input voltage to the FPGA digital port and 1-bit samples obtained by applying the sign-function.

Furthermore, the dither frequency is set to 16 MHz and the power to −4.5 dBm.

In Fig. 11, we report the measured EVM as a function of the average transmitted RF signal power. The error bar at each measurement point marks the standard deviation, estimated based on ten measurements. From the figure, we obtain the following insights: the flatness of the EVM curves in the interval [−60 dBm, −18 dBm] shows that the AGC keeps the average RF signal power constant within 42 dB, which is close to the declared dynamic range of the VGA (45 dB). The VGA has a constant noise figure; hence, the variable gain causes a change in the SNR of the signal. However, we see from the figure that the EVM is approximately constant over the dynamic range of the AGC, which implies that this SNR change is negligible. At the knee appearing just around −18 dBm, the VGA operates at −30 dB gain. For a transmitted power higher than −18 dBm, the VGA operates non-linearly and the signal-to-dither power ratio is sub-optimal. At around −60 dBm, the VGA operates at 15 dB gain. For a transmitted power lower than this, the resulting low SNR, the VGA non-linearity, and the non-ideal dither signal properties contribute to the degradation in EVM.

In the figure, we include also a simulated EVM of 2.7%, obtained by assuming ideal quantization, $\Sigma\Delta$-modulated and lowpass filtered dither signal, and additive white Gaussian noise with power equal to −26.6 dBm prior to dithering. This value of noise power is obtained by assuming a thermal noise of 290 K at the RRH antenna port, and using the noise figures, the gain of the amplifiers, and the bandwidths reported in the component data sheets to calculate the corresponding noise power. Specifically, we use for the LNA, VGA and the RF amplifier the noise figures of 1.5, 12.5 and 2.7 dB, respectively. This noise model is used in the remaining simulations of the receiver. The difference between the simulated and measured EVM within the dynamic range reveals that the noise caused by the amplifiers in the RRH is not the only source of noise in the system. Additional noise may be caused by the optical transceiver components as well as by non-ideal
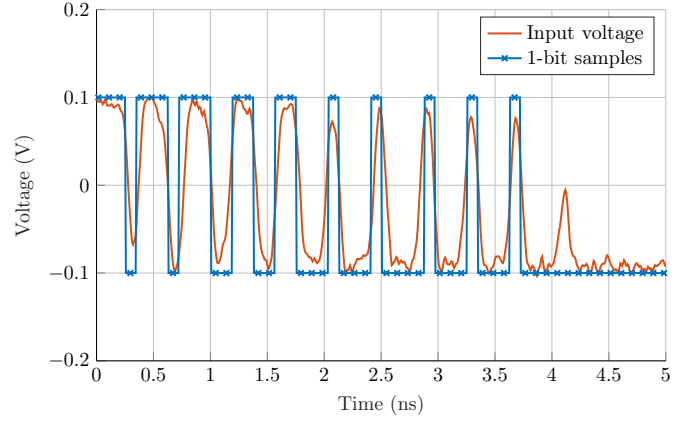
quantization and sampling in the CU. This is discussed further in Section IV-B3.

3) Sampling and Quantization Errors at the CU: As shown in Fig. 11, additional noise sources cause a degradation between 3% and 4% of the EVM. To investigate potential errors introduced during quantization in the SFP+, we connect the output of the SFP+ at the CU to a Keysight *UXR0334A* oscilloscope operating at 128 GS/s, and measure the signal voltage. The result is illustrated in Fig. 12. We start by noting that the input voltage to the FPGA is the output of the comparator at the RRH, after it is passed through the SFP+ at the RRH, the fiber-optical cable, and the SFP+ at the CU. This waveform should ideally be a two-level waveform consisting of pulses of essentially arbitrary width. Indeed, the width of a pulse is determined by the distance between two consecutive zero crossings of the signal obtained by subtracting the dither signal from the filtered RF signal received at the RRH.[2] The actual input voltage is a low-pass filtered version of a two-level waveform. We measure an average rise and fall-time of 54 ps of the input voltage pulses. Hence, if the separation between two zero crossings is less than 54 ps, the input to the FPGA will be significantly distorted. Such a distorted pulse is observed in Fig. 12 at around 4 ns, which may cause the FPGA decision circuitry to generate a bit at random. We assessed the impact of the FPGA sampling by applying in the digital domain the sign function to the recorded input voltage to the FPGA and by using the corresponding quantized signal to perform demodulation. This signal is illustrated by the blue curve in Fig. 12. This resulted in an EVM improvement of about 1%, compared to when the signal is sampled by the FPGA, which partly explains the discrepancy of 3%–4% between simulated and measured EVM.

4) Bandwidth: The bandwidth of the receiver is determined by the bandwidth of the analog components in the RRH. The BPFs have a bandwidth of 100 MHz (2.3–2.4 GHz), and the lower cutoff frequency of the LNA is 2.3 GHz. A reasonable assumption is that the receiver RF-chain in the RRH will not have a flat frequency response

---

[2]Note that the FPGA sampling rate is 10 GS/s. Hence, the minimum duration of the pulses of the 1-bit waveform in Fig. 12 is 0.1 ns.

for input signals whose bandwidth is close to the bandwidth of the RRH components. We investigate the bandwidth limitation of the RRH by presenting the EVM obtainable when receiving signals with symbol rate $R_s \in \{10, 20, 30, 40, 50, 60, 70, 80, 90\}$ MBd. For the case of root-raised cosine pulse shaped single-carrier signal with a roll-off factor of $\alpha = 0.2$, the corresponding bandwidth, computed as $R_s(1 + \alpha)$, is $\{12, 24, 36, 48, 60, 72, 84, 96, 108\}$ MHz, respectively. Since the Rohde & Schwarz *SMU 200A* vector signal generator we used for the measurements in Section IV-B can generate signals with a maximum bandwidth of 50 MHz, we use instead the Agilent Technologies *M9502A* arbitrary waveform generator. The average power of the input signal and of the dither signal are set to −30 dBm and −4.5 dBm, respectively, and we use the dither frequencies $\{16, 22, 31, 37, 45, 61, 69, 77, 81\}$ MHz, which are found via an exhaustive search of the lowest EVM. Furthermore, we compensate for possible inter-symbol interference caused by the hardware by passing the received signal through an L-tap equalizer, before demodulation. The equalizer weights are obtained by least-square channel estimation.

In Fig. 13, we present the measured EVM for single-carrier modulation with single- and ten-tap equalization, and the corresponding simulated EVM.[3] In the simulations, we assume the receiver response to be frequency flat. As shown in the figure, for the scenario considered in this section, the EVM increases as a function of the signal bandwidth. For the case $L = 1$, we see that this increase is roughly linear up to 60 MHz, after which the EVM appears to increase faster than linear. For the case $L = 10$, the linear increase of the EVM holds up to 84 MHz. Furthermore, the resulting EVM values are lower than for the case $L = 1$. This confirms that the frequency response of the RRH is not flat for bandwidths larger than 60 MHz. According to the long-term evolution (LTE) and new radio (NR) standards, the minimum EVM requirement for 16QAM orthogonal frequency-division multiplexing (OFDM) signals is 12.5%, and it is 17.5% for QPSK [25], [26]. Although we use single-carrier signals in this work, we consider an EVM below these levels satisfactory. It follows from the results in Fig. 13 that satisfactory EVM is achieved with the ten-tap equalizer for bandwidths up to 72 MHz for 16QAM, and up to 96 MHz for QPSK.

In Fig. 13 we studied the dependence of the EVM on the signal bandwidth for a fixed sampling rate of 10 GS/s. In Fig. 14, we simulate the impact on the EVM of the sampling rate. As shown in the figure, increasing the sampling rate beyond 10 GS/s is beneficial when the symbol rate is 100 MBd, whereas this has only a modest impact on the EVM when the symbol rate is smaller than 40 MBd. In general, as discussed in [19], the choice of the sampling rate involves a trade-off between system performance and fronthaul rate.

5) OFDM Waveform: So far, all reported measurements pertain the transmission of a single-carrier signal. However, such a signal is not relevant for 5G, which uses instead OFDM. To investigate weather our 1-bit radio-over-fiber fronthaul architecture can support OFDM, we perform an over-cable

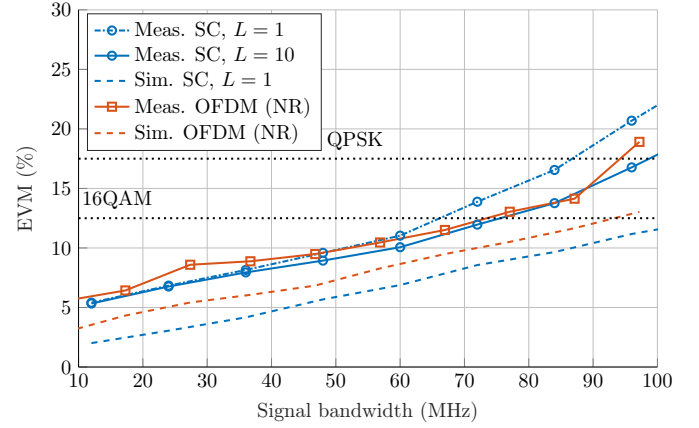[3]The curves with the OFDM label will be discussed in Section IV-B5.

Fig. 13. Average per-symbol EVM for different signal bandwidths. The figure presents measured and simulated results for single-carrier (SC) and OFDM waveforms.
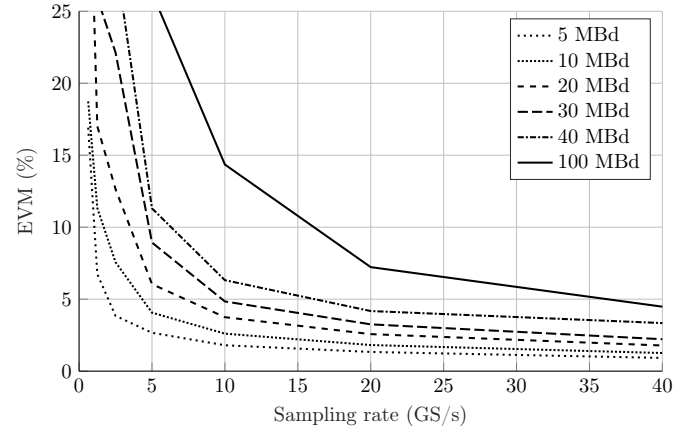

Fig. 14. Simulated average per-symbol EVM for different symbol rates and sampling rates.

measurement in which an OFDM waveform of −30 dBm average power is fed into the antenna port of an RRH.

In accordance with the NR numerology, we use 60 kHz sub-carrier spacing, and a symbol duration (excluding cyclic prefix) of 16.67 μs. Furthermore, we use least-square channel estimation with a delay-domain window of 20 samples on a Zadoff-Chu pilot sequence. The pilot sequence is defined by 3GPP in [27]. The power and frequency of the dither signal are not optimized. Rather, we use the values that were optimal in a similar single-carrier measurement.

In Fig. 13, we present both simulated and measured EVM values as a function of the signal bandwidth. The EVM curves show a similar trend as in the single-carrier case; however, the EVM values are slightly larger. This is expected, since the testbed is not optimized for the OFDM waveform. Despite this, satisfactory EVM is achieved for bandwidths up to 67 MHz for 16QAM and up to 87 MHz for QPSK.

6) Multi-User Reception, In-Band Interference: In practical D-MIMO deployment scenarios, the received signal at the RRH may consist of the superposition of signals from different UEs. The case in which these signals have significantly different power may be problematic for the 1-bit radio-over-

Fig. 15. Simulated and measured average per-symbol EVM versus the SIR, for the case of in-band interference.

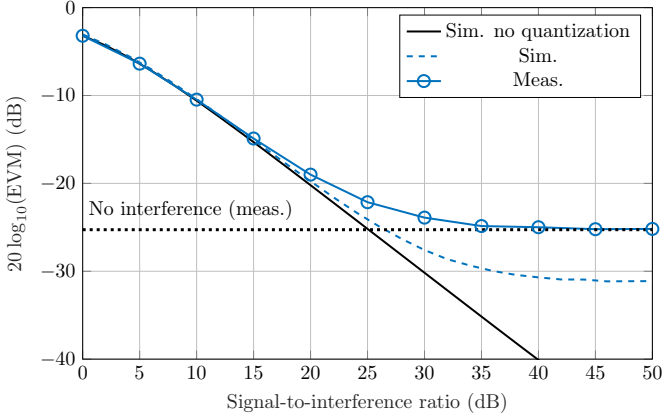
Fig. 16. Simulated and measured EVM versus the SIR, for the case of out-of-band interference.

fiber architecture, since the dither signal will be optimal only for the stronger among the received signals.

To investigate this scenario, we perform an over-cable measurement where the signal connected to the RRH consists of the superposition of two single-carrier signals at the same frequency, one with fixed power $P_{\text{UE1}} = -30$ dBm, and one with variable power P_{UE2}. At the CU, we decode the signal with power P_{UE1}, treating the other signal as noise.

In Fig. 15, we report the measured and simulated EVM in dB as a function of SIR $10\log_{10}(P_{\text{UE1}}/P_{\text{UE2}})$, for a signal with 10 MBd symbol rate. For reference, we include also the simulated EVM for the case of no quantization and infinite SNR, and the measured EVM without interference.

We observe that when the SIR is equal to 15 dB or lower, measured and simulated EVM coincide with the simulated EVM for the case of no quantization and infinite SNR. This implies that, in this regime, the in-band interference dominates over both the thermal and the quantization noise. As the SIR increases, the measured EVM converges rapidly to the one measured in the absence of inter-symbol interference. This suggests that the 1-bit radio-over-fiber fronthaul architecture is no more sensitive to this kind of interference than any other architecture.

7) Multi-User Reception, Out-of-Band Interference: Next, we consider the case of out-of-band interference, i.e., interference from a signal transmitted in a bandwidth adjacent to the one of the desired signal.

We consider the transmission of two 4 MBd signals centered at 2.3475 GHz and 2.3525 GHz, respectively. We fix the power P_{UE1} of the first signal to −30 dBm and vary the power P_{UE2} of the second signal. The dither signal power is optimized for the case in which the SIR is 15 dB. The dither signal frequency is set to 12 MHz and its power to −6 dBm. At the CU, we decode both signals, each time treating the other signal as noise.

In Fig. 16, we present simulated and measured EVM for both signals, as a function of the SIR. As expected, when the SIR is 0 dB, the two signals have the same EVM since they are subject to the same level of interference. As the SIR increases, the EVM of the first signal decreases, although slightly, whereas the EVM of the second signal deteriorates.
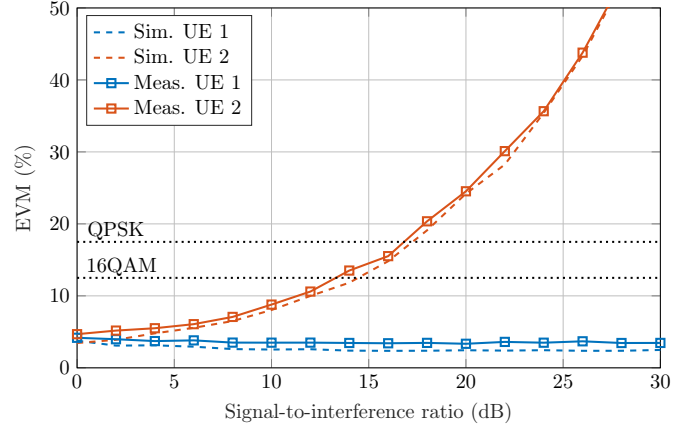
Note though that 16QAM/QPSK can still be supported even when P_{UE2} is 13 dB/16 dB below P_{UE1}, respectively.

In a D-MIMO setting, where the number of RRHs is large compared to the UEs, each UE is likely to generate a strong signal only for a subset of RRHs. By requiring that only these RRHs will serve the UE, one can provide reliable service to all UEs in the network.

V. SYSTEM CALIBRATION FOR OVER-THE-AIR MEASUREMENTS

Having confirmed via point-to-point over-cable measurements that low EVM can be achieved, we now consider over-the-air D-MIMO measurements. In centralized D-MIMO, the uplink channel estimates pertain an effective channel that includes the propagation channel, the RRHs, and the fronthaul. A precoding matrix that is based on the uplink channel estimates will then perform poorly if the uplink and downlink hardware are non-reciprocal, i.e., if they affect the signal amplitude and phase differently. Specifically, the effective baseband uplink channel can be modelled as $\mathbf{H}_{\text{UL}} = \mathbf{R}_{\text{RRH}}\mathbf{H}\mathbf{T}_{\text{UE}}$, where $\mathbf{R}_{\text{RRH}} = \text{diag}(r_1^{\text{RRH}}, \ldots, r_B^{\text{RRH}})$ and $\mathbf{T}_{\text{UE}} = \text{diag}(t_1^{\text{UE}}, \ldots, t_U^{\text{UE}})$ represent the complex-valued gains introduced in the uplink by the B RRHs, the fronthaul, and the U UEs, respectively, and $\mathbf{H} \in \mathbb{C}^{B \times U}$ is the reciprocal channel matrix. Similarly, the effective downlink channel is $\mathbf{H}_{\text{DL}}^{\text{T}} = \mathbf{R}_{\text{UE}}\mathbf{H}^{\text{T}}\mathbf{T}_{\text{RRH}}$, where $\mathbf{R}_{\text{UE}} = \text{diag}(r_1^{\text{UE}}, \ldots, r_U^{\text{UE}})$ and $\mathbf{T}_{\text{RRH}} = \text{diag}(t_1^{\text{RRH}}, \ldots, t_B^{\text{RRH}})$ represent the complex-valued gains introduced in the downlink, by the B RRHs, the fronthaul, and the U UEs.

We assess reciprocity by measuring $\mathbf{H}_{\text{UL}}$ and $\mathbf{H}_{\text{DL}}$, using the setup depicted in Fig. 17, but for a single UE. The vector signal generator acts as a UE transmitter and the oscilloscope acts as a UE receiver. Since this measurement equipment emulates ideal transmitter and receiver hardware at the UE, we assume that $\mathbf{T}_{\text{UE}} = \mathbf{R}_{\text{UE}} = \mathbf{I}_U$. Furthermore, a common clock is shared between the FPGA, the oscilloscope, and the signal generator, to avoid carrier frequency offsets.

We compare the channel estimates $\widehat{\mathbf{H}}_{\text{UL}}$ and $\widehat{\mathbf{H}}_{\text{DL}}$, to assess if reciprocity calibration is required. The estimate $\widehat{\mathbf{H}}_{\text{UL}}$ is obtained by orthogonal pilot transmission from the UE to the RRHs, and $\widehat{\mathbf{H}}_{\text{DL}}$ is obtained by orthogonal pilot transmission
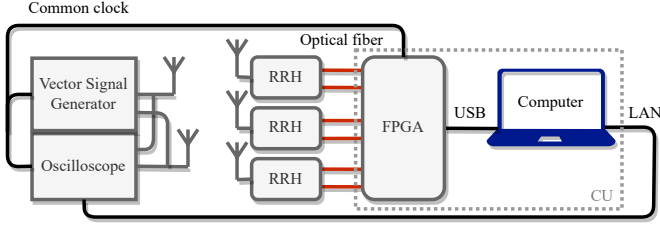
Fig. 17. Setup for multi-user D-MIMO over-the-air measurements. The computer is connected via LAN to the vector signal generator and the oscilloscope, which emulates 2 UE transmitters and receivers, respectively.
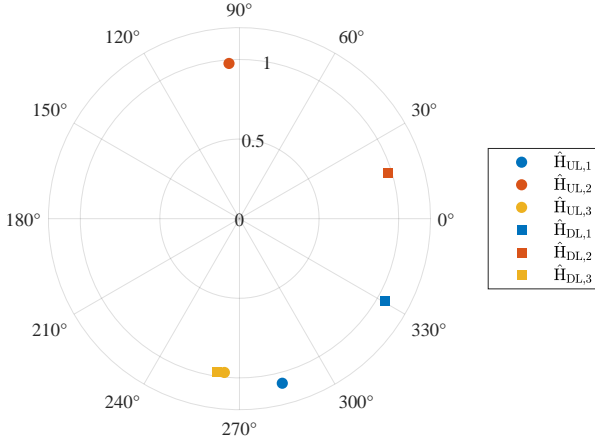

Fig. 18. Channel estimates $\widehat{\mathbf{H}}_{\mathrm{UL}} = \begin{bmatrix} \hat{\mathrm{H}}_{\mathrm{UL},1} & \hat{\mathrm{H}}_{\mathrm{UL},2} & \hat{\mathrm{H}}_{\mathrm{UL},3} \end{bmatrix}^{\mathrm{T}}$ (round marks) and $\widehat{\mathbf{H}}_{\mathrm{DL}} = \begin{bmatrix} \hat{\mathrm{H}}_{\mathrm{DL},1} & \hat{\mathrm{H}}_{\mathrm{DL},2} & \hat{\mathrm{H}}_{\mathrm{DL},3} \end{bmatrix}^{\mathrm{T}}$ (square marks) for three RRHs and one UE.

from the RRHs to the UE. The normalized least-square channel estimates are presented in Fig. 18. We observe that the maximum phase difference between two entries corresponding to the same RRH in $\widehat{\mathbf{H}}_{\mathrm{UL}}$ and $\widehat{\mathbf{H}}_{\mathrm{DL}}$ is $75°$. Hence, the fronthaul is not reciprocal and one needs to perform a reciprocity calibration.

To do so, we use the algorithm proposed in [28], in which calibration is achieved by exchanging signals over the air between all RRHs. The setup is illustrated in Fig. 19. A calibration matrix $\mathbf{C} = \mathbf{T}_{\mathrm{RRH}}\mathbf{R}_{\mathrm{RRH}}^{-1} = \mathrm{diag}(\frac{t_1^{\mathrm{RRH}}}{r_1^{\mathrm{RRH}}}, \dots, \frac{t_B^{\mathrm{RRH}}}{r_B^{\mathrm{RRH}}})$ is found based on an initial guess of $\mathbf{C}$, followed by an iterative gradient search. Specifically, the cost function $J = ||\mathbf{y} - \mathbf{R}_{\mathrm{RRH}}\mathbf{H}\mathbf{T}_{\mathrm{RRH}}\mathbf{x}||^2 + p$, where $\mathbf{y}$ is the received baseband signal, $\mathbf{x}$ is the transmitted baseband signal, and p is a penalty term, is minimized until a threshold tolerance is fulfilled. The precoder based on uplink channel estimates is then multiplied with $\mathbf{C}^{-1}$, to compensate for the lack of reciprocity. We discuss the accuracy of the reciprocity calibration in the next section.

VI. D-MIMO OVER-THE-AIR MEASUREMENTS

In this section, we present multi-user D-MIMO over-the-air measurements for the case in which 3 RRHs serve simultaneously 2 UEs. A vector signal generator and an oscilloscope act as the transmitter and receiver for the two UEs, respectively.
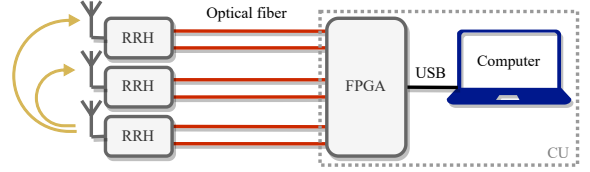

Fig. 19. Setup for over-the-air reciprocity calibration. All RRHs exchange signals with each other.
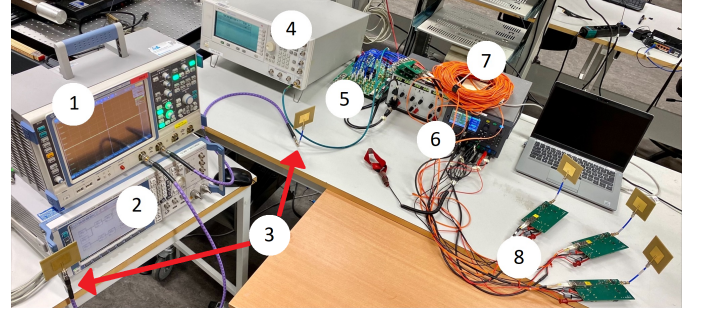

Fig. 20. The testbed used for D-MIMO measurements. The instruments are numbered as 1. Oscilloscope, 2. Vector signal generator, 3. UE antennas, 4. Reference oscillator for FPGA, 5. FPGA evaluation board, 6. DC power supplies, 7. Optical fiber cables, 8. RRHs.

The setup is illustrated in Fig. 17, and a picture of the equipment is presented in Fig. 20. We consider a line-of-sight scenario, with approximately $2\,\mathrm{m}$ distance between the UEs and the RRHs. Furthermore, the two UEs are approximately $1\,\mathrm{m}$ apart and the distance between the RRHs is $1\,\mathrm{m}$. The measurement distances are limited due to the DC power supply that is shared between all RRHs to control the TDD switch.

A. Downlink

We measure the EVM of the received constellation symbols at the UEs when using either $\widehat{\mathbf{H}}_{\mathrm{DL}}$ or $\widehat{\mathbf{H}}_{\mathrm{UL}}$ to create a zero-forcing precoding matrix. Specifically, the zero-forcing precoding matrix based on $\widehat{\mathbf{H}}_{\mathrm{DL}}$ is $\mathbf{P}_{\mathrm{DL}} = \widehat{\mathbf{H}}_{\mathrm{DL}}^*(\widehat{\mathbf{H}}_{\mathrm{DL}}^T\widehat{\mathbf{H}}_{\mathrm{DL}}^*)^{-1}$. Similarly, the zero-forcing precoding matrix based on $\widehat{\mathbf{H}}_{\mathrm{UL}}$ is $\mathbf{P}_{\mathrm{UL}} = \widehat{\mathbf{H}}_{\mathrm{UL}}^*(\widehat{\mathbf{H}}_{\mathrm{UL}}^T\widehat{\mathbf{H}}_{\mathrm{UL}}^*)^{-1}$. We also measure the EVM at the UEs when precoding with either $\mathbf{P}_{\mathrm{UL}}$ or $\mathbf{C}^{-1}\mathbf{P}_{\mathrm{UL}}$ to evaluate the impact of reciprocity calibration.

The constellation diagram and the EVM for the three different precoders are presented in Fig. 21, for the case of a 16QAM signal with $5\,\mathrm{MBd}$ symbol rate. Fig. 21(a) and 21(b) display the received constellation diagrams at the UEs when $\mathbf{P}_{\mathrm{DL}}$ is used for precoding. The measured EVM are 3.7% for UE 1 and 3.8% for UE 2, which implies that the two UEs experience similar SNR and that interference is low. In Fig. 21(c) and 21(d), we present the same results for the case in which $\mathbf{P}_{\mathrm{UL}}$ is used for precoding, without reciprocity calibration. The measured EVM are 89.8% for UE 1 and 84.4% for UE 2, which is clearly not satisfactory. In Fig. 21(e) and 21(f), $\mathbf{C}^{-1}\mathbf{P}_{\mathrm{UL}}$ is used for precoding. The measured EVM are 6.4% for UE 1 and 10.9% for UE 2. This confirms that reciprocity calibration is necessary in our setup for TDD operations based on an uplink pilot transmission.

The degradation in EVM when using $\mathbf{C}^{-1}\mathbf{P}_{\mathrm{UL}}$ (Fig. 21(e) and 21(f)) compared to $\mathbf{P}_{\mathrm{DL}}$ (Fig. 21(a) and 21(b)) for

(a)   (b)

(c)   (d)

(e)   (f)
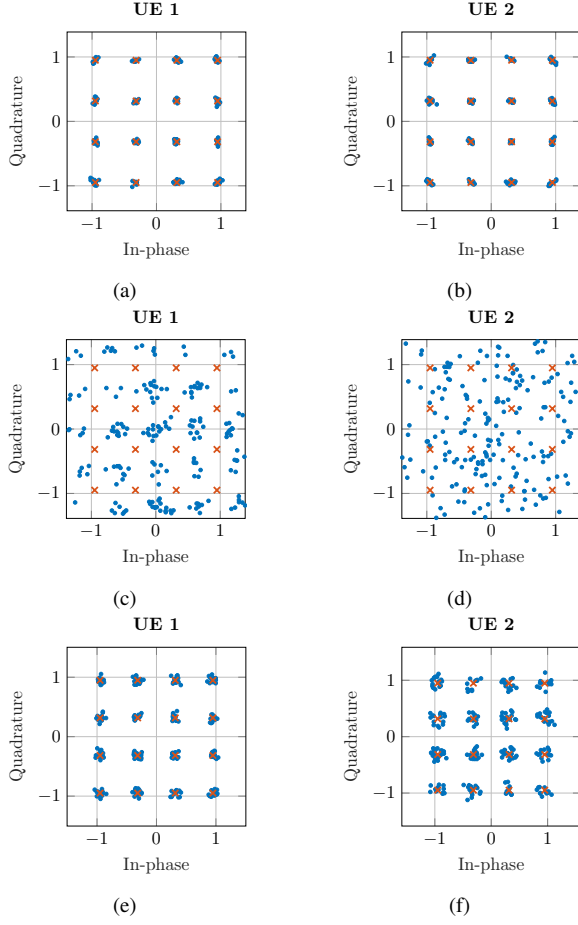
Fig. 21. Constellation diagrams obtained using the precoder based on channel estimation via orthogonal downlink pilots ($\mathbf{P}_{\mathrm{DL}}$), display in (a) 3.7% EVM, and in (b) 3.8% EVM. Constellation diagrams obtained using the precoder based on channel estimation via orthogonal uplink pilots ($\mathbf{P}_{\mathrm{UL}}$), display in (c) 89.8% EVM, and in (d) 84.4% EVM. Constellation diagrams obtained using the precoder based on channel estimation via orthogonal uplink pilots, and calibration is performed ($\mathbf{C}^{-1}\mathbf{P}_{\mathrm{UL}}$), display in (e) 6.4% EVM, and in (f) 10.9% EVM.

downlink precoding indicates imperfect channel estimation and/or calibration. Also, UE 2 experiences larger EVM than UE 1, which indicates that the reciprocity calibration can be improved. A cause of error is that digital signal processing is performed offline, which involves delays in the order of minutes between when the calibration matrix $\mathbf{C}$ is estimated and when it is used.

B. Uplink

We measure the EVM of the received constellation symbols at the CU, after zero-forcing combining. Specifically, 16QAM modulated signals with 5 MBd symbol rate are transmitted from the two UEs simultaneously. The channel estimate $\widehat{\mathbf{H}}_{\mathrm{UL}}$ is obtained from pilot transmission from the UEs. The dither frequency is 12 MHz, and the power is -4.5 dBm. In Fig. 22, the results are displayed in terms of constellation diagrams. The corresponding measured EVM are 5.2% for the signal received from UE 1 and 6.9% for the signal received from UE 2.
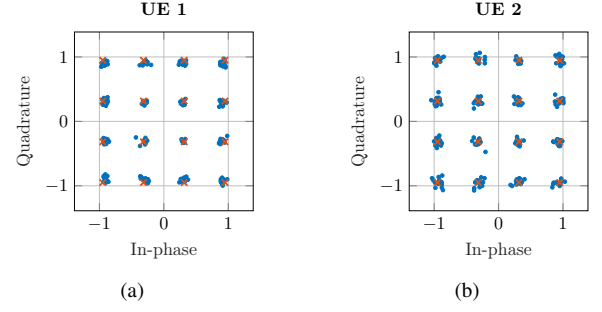

(a)   (b)

Fig. 22. Constellation diagrams for the received signals at the CU displaying in (a) 5.2% EVM and in (b) 6.9% EVM.

We observe that low EVM can be achieved for multi-user uplink transmission, despite the nonlinearity that is introduced by the 1-bit sampling. Furthermore, we observe that, for this particular measurement scenario, the measured EVM per UE is in the same range as the one measured over cable and presented in Fig. 11 (5% to 7%). This confirms that the interference between UEs is low and that quantization-unaware spatial processing provides low EVM per UE.

In summary, we conclude that our testbed supports multi-user D-MIMO when equipped with 3 RRHs serving 2 UEs.

VII. DISCUSSION

The experiments conducted in the paper involve just 3 RRHs. In practice, a much larger number of RRHs need to be considered to reap the benefits of D-MIMO. In the following, we present some remarks on the scalability of the proposed architecture.

- Number of RRHs: The RRHs in the proposed architecture are based on inexpensive, off-the-shelf components. Hence, the proposed architecture remains low-cost even when the number of RRHs increases. Since each RRH needs to be connected to the CU via a fiber-optical cable, the number of RRHs that can be deployed is eventually limited by the number of optical ports available at the CU and by the processing speed of the CU. This makes the proposed architecture promising for local D-MIMO deployments.
- Bandwidth: The sampling rate of 10 GS/s of the proposed architecture effectively limits both the carrier frequency and the bandwidth of the transmitted signal. In the paper, we have assumed a carrier frequency of 2.35 GHz and a bandwidth of up to 100 MHz. Other carrier frequencies and larger bandwidths can be considered, but this would require one to reoptimize the dither signal and to exchange some of the hardware components at the RRHs (specifically, the BPF and the LPF). Also, as shown in Fig. 14, it is beneficial to increase the sampling rate beyond 10 GS/s when the bandwidth of the transmitted signal is larger than 40 MHz. Extensions of the proposed downlink architecture to mm-wave frequencies have been proposed in the literature, and bandwidths up to 700 MHz have been considered [29]. Carrier aggregation, which is used primarily in the downlink, can also be supported; indeed, as shown in [30], $\Sigma\Delta$ modulation is compatible

with carrier aggregation: one just needs to equip the RRHs with a set of reconstruction BPFs at the desired central frequencies.

- Coverage: The 42 dB dynamic range of the AGC at the RRHs determines the coverage of each RRH. If we use the indoor office channel model defined in [31] and we assume that the transmit power of the UE is 23 dBm, we obtain a maximum non-line-of-sight range of 60 m. This implies that each RRH will be able to serve only UEs located within this relatively narrow range. As a result, TDD switching issues are avoided. We would like to emphasize that, using single-mode fiber-optical fronthaul, distances of up to 40 km are supported. Thus, the proposed optical fronthaul architecture is not the limiting factor in terms of coverage.

Measurements for more realistic scenarios, with larger UE–RRH distances, and a larger number of RRHs, will be part of future works.

VIII. CONCLUSION

We have presented both over-cable and over-the-air uplink and downlink measurements for a D-MIMO testbed operating in TDD mode and involving a 1-bit radio-over-fiber fronthaul. In the proposed architecture, frequency up- and down-conversion are performed at the CU. As a consequence, the RRHs do not contain oscillators and mixers, and do not need to be synchronized to enable coherent spatial processing.

We showed that the RRHs can be manufactured using low-cost off-the-shelf components. Our uplink and downlink measured EVM results for the case of 3 RRHs serving 2 UEs show that a satisfactory EVM can be achieved by using linear spatial processing such as zero-forcing, despite the nonlinearity introduced by the 1-bit quantization operation, at the cost of a large oversampling ratio.

We remark that additional RRHs can be added to our measurement setup, with the main limitation being the number of ports available on the FPGA. This makes the proposed architecture scalable.

The distance at which the RRHs can be placed in our testbed is currently limited by the TDD switch synchronization, which is implemented by sharing a DC power supply between the RRHs. As part of future work, we plan to control each switch via a data-channel cable between the CU and the RRH, to be deployed together with the optical fiber, as part of the fronthaul.

REFERENCES

[1] L. Aabel, G. Durisi, I. C. Sezgin, S. Jacobsson, C. Fager, and M. Coldrey, "Distributed massive MIMO via all-digital radio over fiber," in *Proc. Asilomar Conf. Signals, Syst., Comput.*, Pacific Grove, CA, USA, Jun. 2020.

[2] A. Prata, A. S. Oliveira, and N. B. Carvalho, "All-digital flexible uplink remote radio head for C-RAN," in *IEEE MTTS Int. Microw. Symp.*, San Francisco, CA, USA, May 2016, pp. 1–4.

[3] I. C. Sezgin, L. Aabel, S. Jacobsson, G. Durisi, Z. S. He, and C. Fager, "All-digital, radio-over-fiber, communication link architecture for time-division duplex distributed antenna systems," *J. Lightw. Technol.*, vol. 39, no. 9, pp. 2769–2779, May 2021.

[4] S. Maier, W. Kuebart, C. Haslach, U. Seyfried, W. Templ, A. Frotzcher, D. Markert, R. Matz, and A. Pascht, "Class-o base station system with RF pulse-width-modulation in downlink and uplink," in *Asia-Pacific Microw. Conf.*, Melbourne, VIC, Australia, Dec. 2011, pp. 1222–1225.

[5] U. Gustavsson, P. Frenger, C. Fager, T. Eriksson, H. Zirath, F. Dielacher, C. Studer, A. Pärssinen, R. Correia, J. N. Matos, D. Belo, and N. B. Carvalho, "Implementation challenges and opportunities in beyond-5G and 6G communication," *IEEE J. Microwaves.*, vol. 1, no. 1, pp. 86–100, Jan. 2021.

[6] B. M. Khorsandi *et al.*, "Hexa-X architecture for B5G/6G networks," Hexa-X, Tech. Rep., July 2023.

[7] S. Zhou, M. Zhao, X. Xu, J. Wang, and Y. Yao, "Distributed wireless communication system: A new architecture for future public wireless access," *IEEE Commun. Mag.*, vol. 41, no. 3, pp. 108–113, Mar. 2003.

[8] H. Q. Ngo, A. Ashikhmin, H. Yang, E. G. Larsson, and T. L. Marzetta, "Cell-free massive MIMO versus small cells," *IEEE Trans. Commun.*, vol. 16, no. 3, pp. 1834–1850, Mar. 2017.

[9] E. Björnson and L. Sanguinetti, "Making cell-free massive MIMO competitive with MMSE processing and centralized implementation," *IEEE Trans. Wireless Commun.*, vol. 19, no. 1, pp. 77–90, Jan. 2020.

[10] Ö. T. Demir, E. Björnson, and L. Sanguinetti, *Foundations of user-centric cell-free massive MIMO*. Now Publisher, Foundations and Trends in Signal Processing, 2021.

[11] F. Olofsson, L. Aabel, M. Karlsson, and C. Fager, "Comparison of transmitter nonlinearity impairments in externally modulated sigma-delta-over-fiber vs analog radio-over-fiber links," in *Opt. Fiber Commun. Conf. Exhibit. (OFC)*, San Diego, CA, USA, Mar. 2022.

[12] R. Puerta, M. Han, M. Joharifar, R. Schatz, Y.-T. Sun, Y. Fan, A. Djupsjöbacka, G. Maisons, J. Abautret, R. Teissier, L. Zhang, S. Spolitis, M. Wang, V. Bobrovs, S. Lourdudoss, X. Yu, S. Popov, O. Ozolins, and X. Pang, "NR conformance testing of analog radio-over-LWIR FSO fronthaul link for 6G distributed MIMO networks," in *Opt. Fiber Commun. Conf. Exhibit. (OFC)*, San Diego, CA, USA, Mar. 2023.

[13] I. C. Sezgin, T. Eriksson, J. Gustavsson, and C. Fager, "A flexible multi-Gbps transmitter using ultra-high speed sigma-delta-over-fiber," in *IEEE MTTS Int. Microw. Symp.*, Philadelphia, PA, USA, Jun. 2018, pp. 1195–1198.

[14] I. C. Sezgin, M. Dahlgren, T. Eriksson, M. Coldrey, C. Larsson, J. Gustavsson, and C. Fager, "A low-complexity distributed-MIMO testbed based on high-speed sigma-delta-over-fiber," *IEEE Trans. Microw. Theory Techn.*, vol. 67, no. 7, pp. 2861–2872, Jul. 2019.

[15] I. C. Sezgin, T. Eriksson, J. Gustavsson, and C. Fager, "Evaluation of distributed MIMO communication using a low-complexity sigma-delta-over-fiber testbed," in *IEEE MTTS Int. Microw. Symp.*, Boston, MA, USA, Jun. 2019, pp. 754–757.

[16] J. Wang, Z. Jia, L. A. Campos, and C. Knittle, "Delta-sigma modulation for next generation fronthaul interface," *J. Lightw. Technol.*, vol. 37, no. 12, pp. 2838–2850, Jun. 2019.

[17] C. Wu, H. Li, O. Caytan, J. Van Kerrebrouck, L. Breyne, J. Bauwelinck, P. Demeester, and G. Torfs, "Distributed multi-user MIMO transmission using real-time sigma-delta-over-fiber for next generation fronthaul interface," *J. Lightw. Technol.*, vol. 38, no. 4, pp. 4229–4240, Feb. 2020.

[18] S. Jacobsson, L. Aabel, M. Coldrey, I. C. Sezgin, C. Fager, G. Durisi, and C. Studer, "Massive MU-MIMO-OFDM uplink with direct RF-sampling and 1-bit ADCs," in *Proc. IEEE Global Telecommun. Conf. (GLOBECOM)*, Waikoloa, HI, USA, Dec. 2019.

[19] A. Hu, L. Aabel, G. Durisi, S. Jacobsson, M. Coldrey, C. Fager, and C. Studer, "EVM analysis of distributed massive MIMO with 1-bit radio-over-fiber fronthaul," 2023, manuscript submitted for publication.

[20] S. Pavan, R. Schreier, and G. C. Temes, *Understanding Delta-Sigma Data Converters*, 2nd ed., R. Jacob Baker, Ed. Hoboken, NJ, USA: John Wiley & Sons, 2017.

[21] R. F. Cordeiro, A. Prata, A. S. Oliveira, J. M. Vieira, and N. B. De Carvalho, "Agile all-digital RF transceiver implemented in FPGA," *IEEE Trans. Microw. Theory Techn.*, vol. 65, no. 11, pp. 4229–4240, Nov. 2017.

[22] H. du Toit Mouton, B. McGrath, D. Grahame Holmes, and R. H. Wilkinson, "One-dimensional spectral analysis of complex pwm waveforms using superposition," *IEEE Trans. Power Electron.*, vol. 29, no. 12, pp. 6762–6778, Dec. 2014.

[23] Altera Corporation, "Transceiver signal integrity development kit, Stratix V GT edition reference manual," Jan. 2016.

[24] Avago Technologies, "Fiber optical transceiver (AFBR-709SMZ)," Jan. 2013.

[25] 3GPP, "Evolved universal terrestrial radio access (E-UTRA); base station (BS) radio transmission and reception," 2022, TS 36.104 version 17.7.0 Rel. 17.

[26] ——, "NR; base station (BS) radio transmission and reception," 2022, TS 38.104 version 17.7.0 Rel. 17.

[27] ——, "5G; NR; Physical channels and modulation," 2020, TS 38.211 version 16.2.0 Rel. 16.

[28] J. Vieira, F. Rusek, O. Edfors, S. Malkowsky, L. Liu, and F. Tufvesson, "Reciprocity calibration for massive MIMO: Proposal, modeling, and validation," *IEEE Trans. Wireless Commun.*, vol. 16, no. 5, pp. 3042–3056, May 2017.

[29] H. Bao, F. Ponzini, and C. Fager, "Flexible mm-wave sigma-delta-over-fiber MIMO link," *J. Lightw. Technol.*, vol. 41, no. 14, pp. 4734–4742, 2023.

[30] N. V. Silva, A. S. R. Oliveira, U. Gustavsson, and N. B. Carvalho, "A novel all-digital multichannel multimode RF transmitter using delta-sigma modulation," *IEEE Microw. and Wirel. Compon. Lett.*, vol. 12, no. 3, pp. 156–158, Mar. 2012.

[31] 3GPP, "Technical Specification Group Radio Access Network; Study on channel model for frequencies from 0.5 to 100 GHz," 2023, TS 38.109 version 17.1.0 Rel. 17.

Frida Olofsson received the M.Sc degree in engineering physics from the Chalmers University of Technology, Gothenburg, Sweden, in 2020. Since 2020, she has been a Ph.D. candidate at the Department of Microtechnology and Nanoscience. Her research focuses on fiber-optical fronthaul technologies for Distributed MIMO.

Giuseppe Durisi (Senior Member, IEEE) received the Laurea (summa cum laude) and Ph.D. degrees from Politecnico di Torino, Italy, in 2001 and 2006, respectively. From 2002 to 2006, he was with Istituto Superiore Mario Boella, Turin, Italy. From 2006 to 2010, he was a Postdoctoral Researcher with ETH Zurich, Zürich, Switzerland. In 2010, he joined the Chalmers University of Technology, Gothenburg, Sweden, where he is currently a Professor with the Communication Systems Group. His research interests include communication and information theory and machine learning. He was a recipient of the 2013 IEEE ComSoc Best Young Researcher Award for the Europe, Middle East, and Africa region. He is the coauthor of a paper that won the Student Paper Award at the 2012 International Symposium on Information Theory and a paper that won the 2013 IEEE Sweden VT-COM-IT Joint Chapter Best Student Conference Paper Award. From 2011 to 2014, he served as a Publications Editor for the IEEE TRANSACTIONS ON INFORMATION THEORY. From 2015 to 2021, he served as an Associate Editor for the IEEE TRANSACTIONS ON COMMUNICATIONS.

Lise Aabel received the M.Sc. degree in electrical engineering from the Chalmers University of Technology, Gothenburg, Sweden, in 2019. Since 2020, she has been working toward the Ph.D. degree as an Industrial Ph.D. Candidate with Ericsson Research, and the Department of Electrical Engineering, Chalmers University of Technology. Her research focuses on implementation and analysis of distributed antenna systems.

Christian Fager (Senior Member, IEEE) received the Ph.D. degree from the Chalmers University of Technology, Gothenburg, Sweden, in 2003. He is a Full Professor with the Chalmers University of Technology and currently the Head of the Microwave Electronics Laboratory. He has authored or coauthored more than 230 publications in international journals and conferences. His current research interest includes energy efficient and linear transmitters for future wireless communication systems. He is an Associate Editor for IEEE Microwave Magazine and IEEE Microwave and Wireless Technology Letters. He is a member of the Board of Directors in the European Microwave Association (EuMA). He was the recipient of the Chalmers Supervisor of the Year Award in 2018.

Sven Jacobsson received the Ph.D. degree in Electrical Engineering from the Chalmers University of Technology, Gothenburg, Sweden, in 2019. In 2015, he joined Ericsson Research, where he is currently a Senior Researcher working with 5G NR standardization.

Mikael Coldrey received the M.Sc. degree in applied physics and electrical engineering from Linköping University, Linköping, Sweden, in 2000, and the Ph.D. degree in electrical engineering from Chalmers University of Technology, Gothenburg, Sweden, in 2006. He joined Ericsson Research in 2006, where he is currently a Principal Researcher. Between 2017-2020 he was Head of the microwave systems research section at Ericsson Research. Mikael received the Docent (habilitation) promotion from Chalmers in 2020, and was the same year appointed Adjunct Professor in Communication Systems at Chalmers. His main research interests are in the areas of wireless communications, signal processing, advanced antenna systems, and millimeter wave communications for radio access and wireless transport systems.