

Insights on the Uplink Operation of a 1-bit Radio-Over-Fiber Architecture in Multi-User D-MIMO Communication

Lise Aabel, Giuseppe Durisi, *Senior Member, IEEE*, Frida Olofsson, Erik Börjesson, Mikael Coldrey, and Christian Fager, *Fellow, IEEE*

Abstract—We consider a distributed multiple-input multiple-output (D-MIMO) testbed in which, to enable coherent-phase transmission without over-the-air synchronization, the remote radio heads (RRHs) are connected to a central unit via a 1-bit radio-over-fiber fronthaul. Specifically, 1-bit samples of the radio-frequency signal are exchanged over the fronthaul. We investigate via both measurements and simulations based on an accurate model of the testbed hardware, the capability of the proposed architecture to provide uniform quality of services over the coverage area—one of the promises of D-MIMO. Our results are encouraging: for the case in which two user equipments (UEs) communicate over the same 75 MHz signal bandwidth, the measured error-vector magnitude meets the 3GPP New Radio specification of 12.5% for 16QAM across all tested D-MIMO scenarios. We also determine that uplink transmission is a potential bottleneck, due to the limited dynamic range of the automatic gain controller, which prevents the 1-bit quantizer to benefit from dithering. We show that this issue can be mitigated via UE power control.

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

I. INTRODUCTION

DISTRIBUTED multiple-input multiple-output (MIMO) is a wireless network architecture in which multiple, spatially distributed remote radio heads (RRHs) serve user equipments (UEs) cooperatively. The benefits of D-MIMO, compared to conventional centralized architectures, lie in its potential to increase energy efficiency, reliability, and coverage, as well as in its suitability to support integrated communication and sensing operations [1]–[3]. D-MIMO architectures can involve different degrees of cooperation between the RRHs [1], [3]. However, the full benefits of D-MIMO can be released only when the distributed RRHs transmit phase coherent signals to the UEs.

Implementing D-MIMO with phase-coherent transmission requires maintaining radio frequency (RF) phase synchronization among all cooperating RRHs. In conventional RRHs,

digital-to-analog converters (DACs) and analog-to-digital converters (ADCs) operate on baseband (BB) signals, and frequency conversion from BB to RF is performed locally at the RRHs using a mixer and a local oscillator (LO). Unfortunately, LOs are subject to frequency instabilities and phase noise. This yields a time-varying drift of the LOs phases at the RRHs, which needs to be compensated for [4], [5]. This synchronization problem has been studied to a significant extent in the recent literature. We provide next a brief review, which focuses on practical implementations of coherent transmission solutions within the context of time-division duplex (TDD) D-MIMO systems.

A. Synchronization Methods to Enable Coherent Transmission

A prominent line of work considers the conventional RRH design just described, and relies on over-the-air synchronization to provide phase-coherent transmission [4]–[6]. For example, the architecture presented in [6] achieves this through a 802.11-inspired leader-follower approach, in which the follower RRHs continuously estimate and compensate for the phase drift of their LOs. However, as discussed in, e.g., [5], over-the-air synchronization faces unavoidable scalability issues as the number of RRHs grows, for certain topologies.

A drastically different approach is considered in [7]–[9], where the authors focus on less conventional RRH designs and leverage the presence of a wired fronthaul connecting the RRHs to a central unit (CU), to avoid over-the-air synchronization altogether.

In the testbed described in [7], the RRHs are connected to the CU via a fiber-optical fronthaul, over which sigma-delta modulated BB signals are exchanged both in the uplink and the downlink. Specifically, in the downlink, BB signals are sigma-delta modulated at the CU and transmitted over the fronthaul to the RRHs, which perform digital up-conversion using a clock recovered from the data stream. In the uplink, the received RF signals are down-converted in the analog domain using the clock recovered in the downlink, sampled by an ADC, and then sigma-delta modulated and transmitted to the CU. Although this strategy naturally yields synchronization among the RRHs, it results in a complex RRH architecture.

In [8], a mm-wave D-MIMO testbed is presented, in which RRHs and CU are connected via coaxial cables over which intermediate frequency (IF) signals are exchanged, and through which a clock signal is provided to the RRHs. Each RRH

The work of L. Aabel, C. Fager, and G. Durisi was supported in part by the Swedish Foundation for Strategic Research under Grant ID19-0036 and in part by the Sweden's Innovation Agency under Grant 2024-02404.

L. Aabel is with Ericsson AB, 41756 Gothenburg, Sweden, and also with Chalmers University of Technology, 41296 Gothenburg, Sweden (e-mail: lise.aabel@ericsson.com).

M. Coldrey is with Ericsson AB, 41756 Gothenburg, Sweden (e-mail: mikael.coldrey@ericsson.com).

F. Olofsson, G. Durisi, E. Börjesson, and C. Fager are with Chalmers University of Technology, 41296 Gothenburg, Sweden (e-mail: frida.olofsson@chalmers.se; durisi@chalmers.se; erikbor@chalmers.se; christian.fager@chalmers.se).

TABLE I
TDD MULTI-USER D-MIMO TESTBEDS RELYING ON A WIRED FRONTHAUL

Ref.	Fronthaul type	RRHs	UEs	Deployment	Carrier (GHz)	Bandwidth (MHz)
[7]	1-bit BB-over-fiber	4	2	1 m $\times$ 2 m	3.68	46
[8]	IF-over-coaxial	8	8	8 m $\times$ 8 m	28.25	100
[9]	1-bit RF-over-fiber	3	2	1 m $\times$ 2 m	2.35	5
This work	1-bit RF-over-fiber	6	2	3.5 m $\times$ 4 m	2.35	75

then performs up- and down-conversion locally, using the clock signal provided by the CU. Although this ensures phase coherence, the resulting architecture suffers from scalability, due to the short distances over which signals can be exchanged over coaxial cables. Indeed, the coaxial cables are effective only over distances in the order of 10 m. In contrast, a fiber-optic fronthaul can be deployed over several kilometers.

B. The 1-Bit Radio-Over-Fiber Architecture

In [9], we introduced a D-MIMO network architecture upon which this paper expands, which we shall hereon refer to as *1-bit radio-over-fiber architecture*. This architecture involves a CU that is responsible for performing digital frequency up- and down-conversion. Coherent transmission is achieved in this architecture by using only a central LO at the CU and by exchanging 1-bit quantized RF signals between the CU and the RRHs via an optical fronthaul, which enables the use of on-off intensity modulation. A two-level optical signal is generated in the downlink through the use of RF sigma-delta modulation, and in the uplink via RF 1-bit quantization with dithering at the RRHs. While the D-MIMO architectures described in [6]–[8] require delicate clock management—either through optical fiber transmission, coaxial cable connections, or over-the-air synchronization—our architecture circumvents this issue entirely. As an additional advantage, the design of the RRHs is significantly simplified, which is beneficial from a scalability perspective.

1) *Previous Contributions and their Limitations*: The uplink and downlink hardware functionalities of this 1-bit radio-over-fiber architecture have been demonstrated through point-to-point measurements in [9]–[12]. Specifically, in [10], [12] it is demonstrated that satisfactory error-vector magnitude (EVM) performance can be achieved in the downlink using RF sigma-delta modulation, both using a parallel and a serial optical fronthaul configuration. In [9], [11], it is shown that satisfactory EVM performance can be achieved also in the uplink, provided that dithering and oversampling are performed at the RRH. Dithering involves adding to the received signal a suitably designed signal, whose frequency and power are optimized, so as to whiten the quantization noise.

On the negative side, it was also shown in [9] via point-to-point uplink measurements that it is not always possible to operate at the optimal signal-to-dither ratio, because of the limited dynamic range of the automatic gain control (AGC) at the RRH. As a result, whenever the power of the received signal is outside the dynamic range of the AGC, the uplink EVM performance deteriorates significantly.

The multi-user performance of the 1-bit radio-over-fiber architecture is largely unexplored. In [9], because of limi-

tations in the testbed architecture (see Section II-B), multi-user functionalities were demonstrated only for a small-scale deployment scenario involving a coverage area of 1 m $\times$ 2 m in which 3 RRHs serve 2 UEs. Specifically, we showed the feasibility of reciprocity-based coherent downlink transmission over a small bandwidth (5 MHz). Hence, it remains unclear whether the limited dynamic range of the AGC has a negative impact on multi-user performance in more realistic deployment scenarios. Indeed, in the scenario considered in [9], the two UEs are received at a similar power at all RRHs because of their proximity, which makes it possible to guarantee a close-to-optimal signal-to-dither ratio for both UEs signals at each RRH. Although a hardware model for the RRH is presented in [9], this model does not entirely agree with the measurement results and does not capture the limitation in the dynamic range of the AGC. This prevents a simulation-based analysis of its impact on multiuser performance.

2) *Novel Contributions*: In this article, leveraging targeted improvements on the testbed architecture described in [9], we investigate the multi-user performance of the 1-bit radio-over-fiber architecture over a wider bandwidth and within a larger deployment area (see Table I for a comparison between the testbed considered in the present paper and the wired-fronthaul testbeds presented in [7]–[9], in terms of fronthaul type, number of RRHs, number of UEs, size of the deployment area, carrier frequency, and bandwidth). Our specific contributions are as follows:

- We first show that, in the single-UE case, potential performance losses caused by the limited dynamic range of the AGC can be easily avoided, if the RRH are appropriately distributed across the coverage area, and the UE transmit power is properly selected. Under these conditions, satisfactory uplink EVM performance can be achieved uniformly over the entire coverage area.
- Through extensive measurements involving two UEs, placed at difference inter-UE distance within the coverage area, we show that satisfactory EVM uplink performance can also be achieved in the multi-user case. This indicates that the limited dynamic range of the AGC is not a limiting factor for supporting multi-user operation.
- Finally we present an accurate RRH hardware model that effectively captures most hardware impairments and yields EVM predictions closely matching the measured ones. Using this model, we perform simulations to demonstrate how UE power control can be employed to ensure a uniform quality of service.

The rest of this article is organized as follows. In Section II, we describe our improved D-MIMO testbed. Single-user mea-

measurements are described in Section III, whereas multi-user measurements are provided in Section IV. The model of the testbed, together with UE power-control simulations based on this model, are described in Section V. Finally, we provide some concluding remarks in Section VI.

II. THE IMPROVED 1-BIT RADIO-OVER-FIBER TESTBED

We present in this section an overview of the 1-bit radio-over-fiber testbed used to conduct all measurements in this paper. We then highlight the key improvements over the testbed presented in [9], which allow us to consider in this paper multi-user transmission over larger bandwidths as well as a larger deployment scenario.

A. An Overview of the Testbed

We provide in the block diagram depicted in Fig. 1 an overview of the testbed considered in this paper, which involves 6 RRHs serving 2 UEs. In this architecture, the CU consists of a PC, a Xilinx FPGA, and a power supply. Furthermore, the RRHs are equipped with an additional circuit implementing a 12 V to 5 V conversion for the DC power supply and the logic to control the TDD switch, the low-noise amplifier (LNA) gain, and the power amplifier (PA) (see Fig. 2(a)). The PC is responsible for all digital signal processing tasks, including signal modulation and demodulation, dither signal generation, and sigma-delta modulation (SDM).

In the uplink, RF signals are received at the antenna port of each RRH. These signals are amplified by a LNA and then fed to an AGC unit, which ensures that the signal power remains within a prescribed range before reaching one of the two ports of the quantizer. Simultaneously, the CU generates a sigma-delta modulated triangular dither signal and sends it over the downlink optical fronthaul to the RRHs, where it is lowpass-filtered and fed to the other port of the quantizer. The quantizer then generates a binary output signal, based on the sign of the difference of the signals at its two input ports. The quantized RF signal is then converted to the optical domain by the small form-factor pluggable (SFP)28 transmitter circuitry and sent over the optical fiber cable to the CU. The optical signal received at the CU is converted by the QSFP28 to an electrical signal which is sampled by the field-programmable gate array (FPGA) at 25 Gb/s and written to the FPGA memory. We use a personal computer to retrieve the memory content of the FPGA and process the signal, which includes digital down-conversion and demodulation. We emphasize that the AGC has a limited dynamic range, i.e., it cannot always ensure that the RF signal is scaled to the level required for effective dithering. Consequently, when the received RF signal power falls outside the range in which the AGC can adjust its gain, the successive dithering operation becomes ineffective due to a suboptimal signal-to-dither power ratio.

In the downlink, the CU performs signal modulation, digital up-conversion, and RF sigma-delta modulation. The 1-bit output signals from the sigma-delta modulator are converted to the optical domain by the QSFP28 modules and transmitted over individual optical fiber cables, each connecting the CU to a single RRH. At each RRH, an SFP28 module converts

the received two-level optical signal back to the electrical domain. The RF signal is reconstructed by a bandpass filter and provided to a PA. Each RRH is equipped with a patch antenna with 6.5 dBi directivity.

The UE receiver is emulated using a Keysight UXR0334A oscilloscope, while the UE transmitter is emulated using two different signal generators. For single-user measurements, we utilize the Rohde & Schwarz SMCV100B vector signal generator and leverage its remote control options, which facilitate performing repeated measurements. For multi-user measurements, we use the Agilent Technologies arbitrary waveform generator M8190A, because it provides time-synchronized ports that enable the emulation of the transmission from two UEs. Each UE is equipped with a TE Technologies ANT-2.4-CW-HW-SMA omnidirectional half-wave dipole antenna.

B. Targeted Improvements Compared to [9]

As shown in Table I, the measurements conducted with the testbed described in [9]¹ involve signals with only 5 MHz bandwidth. This limitation is due to the relatively low sampling rate at the CU (10 Gb/s) in the previous testbed, which does not provide sufficiently high oversampling rate to support larger bandwidths. In the upgraded testbed, support for much larger bandwidths is enabled by the use of the Virtex UltraScale+ HBM VCU128 FPGA evaluation kit, combined with the optical transceiver modules SFP28 (supporting a rate of 25 Gb/s), and the QSFP28, which operates four parallel channels at 25 Gb/s each (see Fig. 2(b)). Note that the optimal sampling rate to achieve a certain signal quality depends on the deployment scenario, the number of RRHs and the number of active UEs. There is therefore not a direct translation between the sampling rate required to achieve a target EVM for a given bandwidth.

Another limitation of the previous testbed was the small deployment area (1 m × 2 m). With the upgraded testbed, we extended the area to 3.5 m × 4 m by introducing a new solution for distributing the power supply and the logic signals required to control the TDD switch, the LNA gain, and the PA. Specifically, we equipped each RRH with a 12 V-to-5 V converter to ensure a stable 5 V bias voltage, as well as a transistor circuit to maintain proper timing of the logic control signals. Both power and control signals are delivered over the same cable from a power supply located at the CU.

III. SINGLE-USER D-MIMO MEASUREMENTS

One of the key advantages of D-MIMO is its ability to deliver uniform service across the coverage area, provided that a sufficiently large number of RRHs are deployed. In the uplink, this is achieved because, with high probability, the signal from each UE is received at a subset of RRHs with a sufficiently high signal-to-noise-and-interference ratio. As already pointed out, though, in the 1-bit radio-over-fiber architecture considered in this paper, whenever the received RF signal power lies outside the dynamic range of the AGC

¹For clarity, we shall refer to it as *previous* testbed. In contrast, the testbed presented in this paper will be referred to as *upgraded* testbed.






Fig. 1. Block diagram describing the improved testbed architecture.



(a)



(b)

Fig. 2. (a) The RRH equipped with a DC-DC converter and logic control circuit, and (b) the CU connected to QSFP28 modules.

(because it is either too low or too high), performance in terms of EVM is drastically deteriorated because of sub-optimal dithering conditions. Hence, due to this dynamic range limitations, it is unclear whether uniform service across the coverage area can be provided.

The objective of this first investigation is therefore to use the testbed described in Section II to assess service uniformity in terms of EVM, for the case in which the RRHs serve a single UE. The deployment scenario, depicted in Fig. 3, involves furniture, pillars, and cabinets. This yields a rich scattering environment with both line-of-sight and non-line-




Fig. 3. The deployment scenario considered in the measurement campaigns described in Section III and IV. The red circles mark the RRHs, the green circle marks the UEs, and the blue circle marks the CU.

of-sight components.

We evaluate the uplink EVM for the 11 different UE positions illustrated in Fig. 4, where we also depict the positions of the RRHs. To isolate the advantages of a distributed RRH deployment, we report also, for reference, the EVM achieved with our 1-bit radio-over-fiber testbed, for the case in which all RRHs are co-located over one of the tables (see Fig. 4).

The transmitted signal is an orthogonal frequency-division multiplexing (OFDM) waveform with 240 kHz subcarrier spacing over 75 MHz signal bandwidth, and 16-QAM modulation format. The dither signal frequency is 76 MHz and its power is -2 dBm. The UE transmits 4 pilot symbols followed by 1 data symbol. The pilot symbols are used to compensate for frequency offsets and to estimate the channel to each RRH. The channel estimates are then used to perform maximum ratio combining at the CU.

In Fig. 4, we present four heat maps of the measured EVM for the received uplink symbols across all 11 UE positions. We compare two RRHs deployment strategies: distributed and co-located. In the distributed deployment, the RRHs are placed on tables surrounding the UE, whereas in the co-located deployment, all RRHs are arranged in a row at the top-right corner of the area. For both deployments, the UE transmits at two power levels: 0 dBm and -10 dBm.

When the transmit power is 0 dBm, the uplink EVM, averaged over all positions and three consecutive measurements, is $7.0\% \pm 0.8\%$ and $7.2\% \pm 0.5\%$ for the distributed and the co-located deployments, respectively. The similarity in




Fig. 4. Each numbered square represents the measured EVM in the uplink at that position, using our 1-bit radio-over-fiber testbed for both distributed and the co-located deployments over a $3.5\text{ m} \times 4\text{ m}$ area. We consider two UE transmit power levels: 0 dBm and -10 dBm.




Fig. 5. Downlink spectrum at all positions for (a) the distributed deployment and for (b) the co-located deployment.

the EVM values and the small standard deviation across all positions indicate that, in both the distributed and co-located deployments, the received power remains within the dynamic range of the AGC, and an adequate signal-to-dither ratio is maintained. Note that a uniform quality of service can be guaranteed with our 1-bit-over-fiber architecture even in the co-located case, since the AGC, combined with dithering and 1-bit quantization, effectively nullifies the differences in path loss, provided that the received signal power is within the AGC dynamic range.

When we lower the transmit power to -10 dBm, we observe only a small degradation of the average EVM, which is now $8.1\% \pm 0.6\%$, for the distributed deployment case. Note also that the standard deviation remains low. Interestingly, in the co-located case, the degradation is more significant and the standard deviation much larger: the average EVM is $10.5\% \pm 2.2\%$. This implies that, in the co-located scenario, an optimal signal-to-dither ratio cannot be guaranteed at all UE positions due to the low received power.

For completeness, we also verify that uniform quality of service can be guaranteed also in the downlink. We consider the setup in which 2 downlink pilots are transmitted by each RRH in a round robin fashion and used at the UE to estimate the channel. The channel estimates are then fed back to the CU, which uses them to perform maximum ratio transmission.² The transmit power at each RRH is 5 dBm. We report in Fig. 5 the power spectral density of the received signal at the UE across all 11 positions for both distributed and co-located deployments. As shown in the figure, the power spectral density within the bandwidth of interest varies only by 5 dB in the distributed deployment, whereas the variations are more pronounced (15 dB) in the co-located deployment. This confirms that the distributed deployment results in a more uniform power distribution across the coverage area. Interestingly, though, the measured averaged EVM (averaged over all positions and 3 measurements) in the downlink is similar for both deployments: it is $2.9\% \pm 0.3\%$ for the distributed case, whereas it is $2.9\% \pm 0.5\%$ for the co-located case. The similar average EVM in both the distributed and co-located deployments across all positions suggests that the main cause of EVM degradation is not the low received signal-to-noise ratio (SNR), but rather the distortion introduced by the sigma-delta modulator.

It is worth noting that, for both the distributed and the co-located deployments, the average EVM in the uplink is significantly larger than in the downlink. This is due to two main reasons. First, each RRH transmits at an average power of 5 dBm, which is higher than the UE transmit power. Second, the quantization and dithering processes in the uplink introduce more in-band quantization noise compared to the sigma-delta modulation used in the downlink.

IV. MULTI-USER D-MIMO MEASUREMENTS

After having verified that uniform quality of service over the coverage area can be guaranteed in the case of single-user transmission, we next extend our analysis to the multi-user case. Our main aim is to determine how and in which circumstances the limitation in the dynamic range of the RRHs affect multi-user operations. We consider the same distributed and co-located deployment scenarios described in Section III (again, the co-located scenario is considered to disentangle the effects caused by distributing the RRHs from the ones caused by the signal distortions introduced by the 1-bit distributed-over-fiber architecture), but assume this time that 2 UEs are active within the coverage area. We use zero-forcing precoding for downlink transmission and zero-forcing combining in the uplink. The transmitted power from each UE is -5 dBm. All other parameters are as detailed in Section III.

As far as the UEs positions are concerned, we consider the three different configurations illustrated in Fig. 6(a).

Configuration 1: The UEs are placed 0.8m apart in the center of the coverage area. This arrangement provides both UEs with similar channel gains to all RRHs in both the distributed and the co-located deployment.

Configuration 2: The UEs are located in opposite corners of the coverage area. In the distributed deployment, each UE is spatially close to a different set of RRHs. In the co-located deployment, one UE is significantly closer to the RRHs than the other, leading to a significant disparity in the received signal power.

Configuration 3: The UEs are placed 10cm apart in the center of the coverage area. This proximity increases the probability of the UEs experiencing highly correlated wireless channels, which makes exploiting spatial multiplexing challenging. With this configuration, we test the capability of our 1-bit radio-over-fiber architecture to separate the two UEs in the distributed and co-located deployment cases.

Configurations 2 and 3 were chosen to analyze stress-test scenarios, in which we expect performance degradation in at least one of the deployment scenarios, whereas configuration 1 was chosen to explore a scenario in which it is reasonable to expect similar performance in both deployment scenarios.

We report in Fig. 6(b) and Fig. 6(c) the measured EVM, averaged over three consecutive measurements as well as the corresponding standard deviation, in the uplink and the downlink for both UEs, all three configurations, and both deployment scenarios. As shown in the figure, in the distributed deployment scenario, a relatively uniform EVM can be achieved both in the uplink and in the downlink for all three configurations. Furthermore, this EVM is below the level required by 3GPP for 16QAM transmission. On the contrary, in the co-located case, we observe that the EVM changes drastically across configurations and exceeds the 3GPP requirement in the uplink for at least one UE. When comparing the EVM between the two deployments, we observe that UE 2 experiences increased EVM in the co-located uplink scenario. This is due to sub-optimal dithering conditions caused by the stronger power received at all RRHs from UE 1. This effect is even more pronounced in Configuration 2, because the received power from UE 2 is considerably lower. Additionally,

²An alternative to this approach is to perform reciprocity calibration and downlink multi-user transmission based on uplink channel estimation. In [9], we have demonstrated that this alternative approach is feasible with our 1-bit radio-over-fiber testbed.




Fig. 6. (a) Positions of the two UEs for the three configurations considered in Section IV; (b) average measured EVM and standard deviation for both distributed deployment and (c) co-located deployment. The dashed lines marks the 3GPP EVM requirement for 16QAM transmission.

a notable difference in downlink performance between the two UEs is evident in Configuration 2, indicating that the channel matrix is ill conditioned. In Configuration 3, both the uplink and downlink EVM are large in the co-located case, and similar for both UEs, suggesting that, different from the distributed case, spatial multiplexing is not possible.

To summarize, our measurements confirm that the reduced dynamic range of our RRHs is more deleterious in terms of performance in the multi-user case than in the single-user case, as illustrated by the unsatisfactory performance achievable in the co-located deployment. Encouragingly, though, such limited dynamic range does not seem to affect significantly the performance in the distributed deployment case. It is also interesting to note that the difference between the EVM achieved in the uplink and in the downlink in the multi-user case is not as significant as in the single-user case, which suggests that the downlink performance is now limited by multi-user interference. We expect the conclusions drawn from this measurement campaign to generalize to variations in user configurations.

V. AN ACCURATE HARDWARE MODEL FOR OUR TESTBED

The measurements reported in Section IV using our 1-bit radio-over-fiber testbed involve only six RRHs and two UEs. In order to investigate larger scenarios, but also explore a larger range of transmit signal power values, we present in this section an accurate hardware model for the testbed, with the specific aim to investigate uplink performance and the impact of UE power control. This model is obtained by accounting for the noise figure associated with the amplifiers, the limited dynamic range of the AGC, and the imperfections in the pulse generated by the optical transceivers.

TABLE II
PARAMETER VALUES FOR THE RECEIVER MODEL PRESENTED IN FIG. 6.

Component identifier	B	T_e	G
LNA	400 MHz	119 K	24 dB
VGA	3 GHz	4867 K	See eq. (1)
DA	180 MHz	319 K	15 dB
LPF	180 MHz	—	—
BPF	100 MHz	—	—

A. A Model for the Uplink Fronthaul Signal

Our hardware model is based on an accurate characterization of the signal exchanged over the fronthaul in the uplink. Such characterization is based on the block diagram presented in Fig. 7. There, we describe how the inputs to each RRH, i.e., the received RF signal y and the sigma-delta-modulated dither signal d from the CU are processed at the RRH to obtain the two-level waveform z at its output. Both y and d pass through amplifiers, which generate noise and degrade the SNR. We model the power of the additive noise added by each amplifier as indicated in Fig. 7, where k is the Boltzmann constant, B is the component-specific bandwidth and $T_e = (F - 1)T_0$ is the equivalent noise temperature. Here, F denotes the noise figure [13, Ch. 3.5] and $T_0 = 290$ K. The values of B , T_e and of the gain G of each component are given in Table II.

We next detail the modeling of two critical components: the AGC and the quantizer.

AGC: The AGC at the RRH consists of a variable gain amplifier (VGA), from which a small portion of the output signal is tapped and fed to a logarithmic power detector that generates the output voltage V_{GAIN} . This output voltage determines the gain of the VGA, which is thus a function of




Fig. 7. Block diagram of the proposed RRH model when the RRH operates in receive mode.




Fig. 8. Measured AGC control voltage V_{GAIN} and the corresponding modeled amplifier gain.

its input signal power. The logarithmic nature of the power detector ensures that V_{GAIN} is a linear function of the input power in dB. To construct a model of the VGA, we measure its control voltage V_{GAIN} for different received power levels, measured at the antenna port of the RRH. We then map this voltage to a gain value according to the component datasheet [14]. In Fig. 8, we compared the measurement result with a linear model of V_{GAIN} as a function of the received power in dB. We observe from the figure that V_{GAIN} is indeed approximately a linear function of the input power in dB, within the dynamic range of the AGC, and that a linear model fits the measured gain control voltage, apart from a small deviation at high received power values. According to the component datasheet, the gain in dB of the VGA is directly proportional to V_{GAIN} . We model therefore the gain of the VGA as

$$G_{\text{VGA}}(P_{y_1}) = \begin{cases} 10 \text{ dB}, & \text{if } P_{y_1} < -41 \text{ dBm} \\ -35 \text{ dB}, & \text{if } P_{y_1} > 4 \text{ dBm} \\ -P_{y_1} - 31 \text{ dBm}, & \text{otherwise,} \end{cases} \quad (1)$$

where P_{y_1} is the average power of the signal y_1 , measured in dBm. Note that, when the received power drives the VGA beyond its dynamic range, the gain is modeled as a constant.

Quantizer: As already pointed out, we use the SFP28 module at the RRH to generate a two-level pulse train, which is obtained by tracking the sign of the difference between the filtered and amplified RF signal y_2 and the filtered and

amplified dither signal d_2 . It is important to note that the transition between the two levels in the pulse train is not instantaneous. Rather, it is constrained by the rise and fall time of the module. This pulse train is converted to the electrical domain by the QSFP28 at the CU and then sampled at the digital ports of the FPGA. The distortions caused by the non-instantaneous rise and fall times imply that some of the pulses have an amplitude that is too low to be detected by the circuitry of the FPGA. This results in missed pulses during the sampling operation.

To investigate this phenomenon, we compare the bit sequence sampled by the FPGA with the one obtained by sampling the same signal using an oscilloscope and quantizing it with the sigum function. Since the oscilloscope has a high amplitude resolution, the output of the sigum function approximates well that of an ideal 1-bit converter. Our investigation reveals that pulses are sampled correctly by the FPGA only when the differential input signal to the SFP28 is larger than the threshold voltage $V_{\text{th}} = 10 \text{ mV}$. To model this behaviour, we assume that, whenever the magnitude of the differential input is below V_{th} , the corresponding bit is drawn independently at random from a Bernoulli distribution with parameter $1/2$ (denoted as $\text{Bern}(1/2)$):

$$z = \begin{cases} 1, & y_2 - d_2 > V_{\text{th}} \\ 0, & y_2 - d_2 < -V_{\text{th}} \\ \sim \text{Bern}(1/2), & |y_2 - d_2| < V_{\text{th}}. \end{cases} \quad (2)$$

Validation of the resulting model: To validate the accuracy of the proposed model, we conduct point-to-point measurements of the EVM for the case in which a signal transmitted at different power levels is sent from a signal generator to the antenna port of each one of the six RRHs via a cable. The measurement is conducted for both 5 MHz and 75 MHz signal bandwidth, to identify potential bandwidth-dependent distortions. The measured EVM is then compared to the simulated EVM for the same signal power levels. The results are presented in Fig. 9, where dashed lines represent the simulation results, solid lines represent the average measured EVM, averaged over all 6 RRHs, and the shaded areas denote the range between maximum and minimum measured EVM at each point. As shown in the figure, the simulated EVM curves closely match the minimum measured EVM at each point for both bandwidths, indicating that the hardware model effectively captures the noise characteristics and dynamic range of the uplink hardware in both cases.




Fig. 9. The measured and simulated EVM per RRH as a function of the received power. Solid lines represent the average EVM across all RRHs, and shaded areas represent the range between maximum and minimum EVM at each point.

The discrepancies in the measured EVM across the RRHs are probably due to tuning errors when manually adjusting the output voltage from the AGC, as well as hardware irregularities caused by the soldering process. In a few measurements at very low and very high power levels, the reported EVM is lower than the simulated value, likely due to nonlinear effects occurring when the VGA operates outside its dynamic range.

B. Comparison with Infinite-Precision Architecture

As illustrated in Fig. 6(c) for the colocated case, the EVM performance of our 1-bit radio-over-fiber architecture is unsatisfactory whenever there is a large difference between the signal power received by all RRHs from each of the UE. In this section, we use the model for the uplink fronthaul signal presented in Section V-A to shed more light on this phenomenon. In particular, we show that it may occur also in the distributed case, and that the corresponding EVM degradation is particularly pronounced with our 1-bit radio-over-fiber architecture, whereas it would be much milder if an undistorted version of the RF signal at each RRH would be made available at the CU.³

Specifically, we consider the scenario depicted in Fig. 10(a). In this scenario, UE 1 is in close proximity of all three RRHs, whereas UE 2 is further away. This setup ensures that the received power contribution from UE 2 is weaker than that from UE 1 at all RRHs. To model the propagation channel, we use the line-of-sight path loss model for the urban micro scenario given in [15]. No small-scale fading is considered, for simplicity. The other system parameters are the same as in Section IV.

In Fig. 10(b), we report the EVM simulated using the model presented in Section V-A. For reference, we also report the EVM achievable for the case in which the undistorted RF signal is available at the CU, which we denote by ∞ -bit (infinite-resolution converters). As expected, in both cases UE 2 experiences a worse EVM than UE 1; however, while

³This would be difficult to achieve, due to unavoidable distortions in the transmission over the optical fiber.




Fig. 10. (a) Simulation scenario involving two UEs and three RRHs (b) the EVM in the uplink achieved using 1-bit conversion and dithering or an infinite-resolution converter.

the EVM levels are still well below the 3GPP requirements in the ∞ -bit case, in the 1-bit case these values are significantly exceeded by UE 2. This confirms that our 1-bit radio-over-fiber architecture is more sensitive to receive-power imbalances in the uplink compared to a conventional architecture that uses high-precision converters.

C. The Benefits of UE Power Control

One way to address the EVM degradation illustrated in Fig. 10(b), is to implement power control. Each RRH sends at regular intervals beacons that are used by all UE to estimate the received power. Then this information is conveyed to the CU, that decides the power at which each UE should transmit (and typically, the subset of RRHs that should serve each UE, according to the user-centric-design principle [16]). This UE power control strategy has also the additional advantage in our 1-bit radio-over-fiber architecture to prevent that the AGC operates outside its dynamic range (see Fig. 9).

To investigate the benefits of UE power control on the per-UE performance in terms of EVM, we consider the scenario depicted in Fig. 11(a), which involves 5 UEs and 12 RRHs. As shown in the figure, UE 5, which is located at the center

of the coverage area, has the greatest distance to its nearest RRH among all UEs. UEs 1–4 are arranged in a symmetric grid, which ensures that they experience similar transmission conditions. Consequently, if all UEs transmit at the same power level, we expect UE 5 to exhibit the highest EVM. To confirm this intuition and explore the effect of power control, we plot in Fig. 11(b) the simulated EVM for all UEs, as a function of the transmitted power of UEs 1–4, which is assumed to be identical. In our simulation, UE 5 transmits at a constant power of 10 dBm. The figure illustrates that when all UEs transmit at 10 dBm (no power control), UE 5 exhibits, as expected, the highest EVM value. As the transmitted power from UEs 1–4 is reduced, the EVM for UE 5 improves. In particular, all UEs experience the same EVM when the transmit power from UE 1–4 is about 9 dB. A further reduction of the transmit power from UEs 1–4 is initially beneficial for UE 5. However, when the power falls below 6 dB, the EVM of UE 5 starts to degrade again. This is because the channel estimates for UE 1–4 deteriorate, which impact significantly the performance of the zero-forcing combiner.

D. Towards a Power Control Algorithm

We conclude this section by outlining some preliminary ideas for a potential power control algorithm. A common design objective for such algorithms is to ensure fairness among users. This is typically achieved through a max-min fairness criterion, which seeks to maximize the minimum user spectral efficiency in the system [16], [17]. As shown in [16], for D-MIMO systems with infinite-precision fronthaul, closed-form expressions for a lower bound on each users' achievable rate can be derived, and then optimized for max-min fairness using a fixed-point algorithm.

However, these closed-form expressions do not apply to the 1-bit radio-over-fiber architecture considered in this paper, due to the nonlinearity introduced by the 1-bit quantizer. Although analyzing the spectral efficiency achievable with this architecture lies beyond the scope of this paper, we briefly outline how such an analysis could be conducted and the challenges it would pose for power-control design. One possible approach, adopted in [18] for theoretical EVM characterization, is to linearize the quantizer using Busgang's theorem [19] under the simplifying assumption that the transmitted signal and the dither signal are Gaussian distributed. We believe that this approach could yield a spectral efficiency lower bound that explicitly accounts for the quantization error and for the presence of the dither signal, with the latter modelled as a function of the received power and the VGAs gain (1). The next step would be to determine whether the commonly used fixed-point optimization algorithm can be applied to the resulting spectral-efficiency expression.

VI. CONCLUSION

We have presented hardware improvements on the 1-bit radio-over-fiber testbed introduced in [9], aimed at increasing both the bandwidth and the size of the deployment scenario. Thanks to these improvements, we were able to conduct measurement campaigns aimed at assessing whether one can




Fig. 11. (a) Simulation scenario involving 5 UEs and 12 RRHs and (b) the per-user EVM in the uplink as a function of the transmitted power from UE 1 through 4.

achieve the uniform quality of service promised by D-MIMO with our 1-bit radio-over-fiber architecture. Our initial results are promising (see Fig. 6(b)). We have also identified that uplink transmission is a potential bottleneck, due to the limited dynamic range of the AGC (see Fig. 9), which may prevent the 1-bit quantizer from operating at the optimal signal-to-dither power ratio. As illustrated in Fig. 6(c), this issue yields a notable performance degradation when there is a significant difference in the power received from the UEs at all RRHs. By leveraging the accurate simulation model introduced in Section V-A, we have shown that this issue can be mitigated via UE power control. These findings highlight the need to devise novel UE power control and UE-RRH association strategies that explicit account for the specific hardware limitations imposed by the 1-bit radio-over-fiber architecture considered in this paper.

REFERENCES

- [1] 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.
- [2] M. F. De Guzman *et al.*, "Final results of 6G radio key enablers," Hexa-X-II, Tech. Rep., Feb 2025.

- [3] H. Guo, H. Wymeersch, B. Makki, H. Chen, Y. Wu, G. Durisi, M. F. Keskin, M. M. Moghaddam, C. Madapatha, H. Yu, P. Hammaberg, H. Kim, and T. Svensson, "Integrated communication, localization and sensing in 6G D-MIMO networks," *IEEE Wireless Commun. Mag.*, vol. 32, no. 2, pp. 214–221, Feb. 2025, accepted for publication.
- [4] N. Shandi, J. M. Merlo, and J. A. Nanzer, "Decentralized picosecond synchronization for distributed wireless systems," *IEEE Trans. Commun.*, vol. 73, no. 6, pp. 4425–4438, Jul. 2024.
- [5] E. G. Larsson, "Massive synchrony in distributed antenna systems," *IEEE Trans. Signal Process.*, vol. 72, pp. 855–866, Jan. 2024.
- [6] E. Hamed, H. Rahul, M. Abdelghany, and D. Katabi, "Real-time distributed MIMO systems," in *SIGCOMM 2016 - Proceedings of the 2016 ACM Conference on Special Interest Group on Data Communication*, Florianópolis, Brazil, August 2016.
- [7] A. Vandierendonck, K.-L. Chiu, C. Meysmans, F. Zardosht, X. Wang, H. Li, P. Demeester, and G. Torfs, "A bi-directional distributed multi-user MIMO testbed using digital sigma-delta-over-fiber," *J. Lightw. Technol.*, vol. 43, no. 4, pp. 1595–1603, 2024.
- [8] T. Kaneko, T. Kuwabara, N. Tawa, and Y. Maruta, "Bit efficiency of distributed- and collocated-massive MIMO base station systems in OTA measurement and simulation," *Intern. J. Microw. Wireless Technol.*, 2024.
- [9] L. Aabel, S. Jacobsson, M. Coldrey, F. Olofsson, G. Durisi, and C. Fager, "A TDD distributed MIMO testbed using a 1-bit radio-over-fiber fronthaul architecture," *IEEE Trans. Microw. Theory Techn.*, vol. 72, no. 10, pp. 6140–6152, Oct. 2024.
- [10] F. Olofsson, L. Aabel, T. Eriksson, M. Karlsson, and C. Fager, "Sigma-delta-over-fiber with WDM serial connection for distributed MIMO," *J. Lightw. Technol.*, vol. 43, no. 5, pp. 2151–2162, 2025.
- [11] 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.
- [12] 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.
- [13] D. M. Pozar, *Microwave and RF wireless systems*. John Wiley & Sons, 2001.
- [14] Analog Devices, "ADL5330," <https://www.analog.com/en/products/adl5330.html>, accessed: 2025-06-19.
- [15] 3GPP, "Technical Specification Group Radio Access Network; Study on channel model for frequencies from 0.5 to 100 GHz," 2024, TS 38.901 version 18.0.0 Rel. 18.
- [16] 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.
- [17] 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.
- [18] A. Hu, L. Aabel, G. Durisi, S. Jacobsson, M. Coldrey, and C. Fager, "EVM analysis of distributed massive MIMO with 1-bit radio-over-fiber fronthaul," *IEEE Trans. Commun.*, vol. 72, no. 11, pp. 7342–7356, Nov. 2024.
- [19] J. J. Bussgang, "Crosscorrelation functions of amplitude-distorted Gaussian signals," Res. Lab. Elec., Cambridge, MA, Tech. Rep. 216, Mar. 1952.