

Wideband Synchronized Bidirectional Distributed MU-MIMO With Sigma-Delta RoF: E2E Experimental Gains Compared to Collocated MIMO

MUTEEN MUNAWAR^{ID} (Student Member, IEEE), ACHIM VANDIERENDONCK (Student Member, IEEE),
MAMOUN GUENACH^{ID} (Senior Member, IEEE), CARO MEYSMANS (Student Member, IEEE),
PARADISE ZARDOSHT, GUY TORFS (Senior Member, IEEE), AND INGRID MOERMAN^{ID} (Fellow, IEEE)

Interuniversity Microelectronics Centre (IMEC), Leuven 3001, Belgium
Internet Technology and Data Science Laboratory (IDLab), IMEC, Ghent University, Ghent 9050, Belgium

CORRESPONDING AUTHOR: M. MUNAWAR (muteen.munawar@imec.be)

This work was supported in part by the Research Foundation Flanders (FWO) under Grant A2582_00_01_01, in part by the Project Strip-Link Multiple-Input Multiple-Output (MIMO) under Grant FWO-97707, and in part by the Interuniversity Microelectronics Centre (IMEC) Advanced Radio Frequency (ARF) Program under Grant 97700/01345.

This article has supplementary downloadable material available at <https://doi.org/10.1109/OJCOMS.2026.3706506>, provided by the authors.

ABSTRACT Cell-free (CF) massive MIMO (mMIMO), also referred to as distributed MIMO (dMIMO) in indoor environments, has emerged as a key enabling technology for beyond-5G wireless systems due to its macro-diversity, favorable propagation, and joint multi-access-point (AP) processing gains. While extensive theoretical studies have demonstrated its potential, experimental validation remains limited. In this work, we present an end-to-end, multi-user, multi-stream, wideband OFDM experimental dMIMO platform, consisting of a centralized processing unit connected via optical fiber fronthaul links to geographically distributed APs. Tight synchronization among APs is achieved using sigma-delta modulated radio-over-fiber (SDRoF), combining the synchronization advantages of analog RoF with the cost efficiency of digital RoF. Using this platform, we experimentally evaluate three CF processing architectures, centralized (L4), hybrid centralized-distributed (L3), and fully distributed (L2), benchmarked against a collocated cellular MIMO system, within a complete end-to-end physical-layer framework comprising packet structure design, pilot signaling, wideband OFDM processing, and spatial processing algorithms. We discuss the uplink (UL) and downlink (DL) timing and frequency synchronization methods, as well as packet demultiplexing, in detail. Experimental results demonstrate substantial macro-diversity and favorable propagation gains. For an 8×4 MIMO scenario with four collocated users and two APs separated by 3 m, centralized CF processing (L4) achieves a minimum per-user signal-to-interference-plus-noise-plus-distortion ratio (SINDR) of 24 dB, versus only 8 dB for the collocated cellular configuration ($\lambda/2$ AP spacing), yielding a 16 dB gain. Under shadowing, the CF architecture maintains 30 dB minimum SINDR compared to 22 dB for the cellular system at $K = 2$ users. Furthermore, we experimentally validate UL-DL duality and quantify the reciprocity mismatch impact, demonstrating a 10 dB DL degradation without hardware calibration in a single-user (8×1) configuration. To the best of our knowledge, this work represents one of the first comprehensive experimental validations of wideband dMIMO and quantifies its practical gains over traditional cellular architectures.

INDEX TERMS Distributed MIMO, collocated MIMO, wideband OFDM, multi-user multi-stream, spatial processing, shadowing, favorable propagation, radio-over-fiber, sigma-delta modulation.

I. INTRODUCTION

TO MEET the high data-rate and reliability demands of beyond-5G wireless networks, cell-free (CF) massive multiple-input multiple-output (mMIMO) is envisioned as

one of the key enabling technologies [1]. By definition, CF refers to a wireless network that does not rely on the traditional cellular concept, in which each cell uses a separate frequency spectrum and communicates only with users within its boundary. In contrast, as the name implies, CF systems

operate without fixed cell boundaries, allowing the entire network to share a common frequency, time, and spatial resource pool [2]. Furthermore, CF systems are distributed by construction: a large number of access points (APs) are deployed across the region of interest.

Such a network configuration, which distributes APs across the region while sharing frequency, time, and spatial resources, provides two key benefits in wireless communication [3]:

- 1) *Macro-diversity*: this captures the large-scale fading benefit of CF mMIMO. When multiple APs jointly serve a single user, that user is less likely to experience complete shadowing because the APs are distributed across the coverage area, so it is highly unlikely that all of them are simultaneously blocked.
- 2) *Favorable propagation*: this captures the small-scale fading benefit. Even when two user equipments (UEs) are geographically collocated, APs at different locations can jointly orthogonalize the UEs' channels through simple linear processing, because each AP observes the UEs from a distinct spatial angle.

A. RATIONALE FOR CF AND IMPLEMENTATION CHALLENGES

The macro-diversity and favorable-propagation benefits ensure reliable, uniform UE performance throughout the coverage region, as established in numerical studies of CF mMIMO [4]. Realizing them in practice raises two classes of challenges: signal-processing and synchronization.

Signal processing. Two paradigms exist [3]: *centralized processing*, where the CPU has global CSI and optimizes beamforming and resource allocation while the APs act as radio units; and *distributed processing*, where the APs compute precoders locally from their own CSI and the CPU optimizes only statistical parameters. Centralized processing maximizes performance at the cost of CPU complexity; distributed processing offloads computation to the APs but degrades in dense UE scenarios [5]. MMSE beamforming with convex-optimization-based power allocation [6], [7], [8], [9] is well-suited to numerical studies but impractical for real-time operation given indoor coherence times of a few milliseconds [10], motivating experimental validation of real-time-implementable spatial-processing algorithms.

Synchronization. Multiple geographically distributed APs connected to a CPU via fronthaul cables must remain tightly synchronized for the coherent-transmission benefits to materialize [5]. Most numerical CF studies assume perfect synchronization [1] or use simplified error models [11], [12]; achieving sub-10 ps real-time synchronization across many APs remains a significant practical challenge [13].

B. THEORETICAL LITERATURE

The CF mMIMO architecture, in which a large number of APs jointly serve all UEs, was first explicitly defined and analyzed in [4]. The L4 (centralized) and L3 (hybrid) processing tiers were introduced in [3], and a scalable

variant via dynamic local clustering was proposed in [14]. Improved distributed processing through iterative refinement was studied in [15], and additional numerical studies have addressed constrained fronthaul bandwidth [16], [17], low-resolution ADCs [18], phase noise [11], imperfect CSI [19], functional-split trade-offs [5], multi-stream fairness [6], and asynchronous synchronization [12], [13], [20]. More recently, joint AP-UE association together with precoding has been studied for stacked-intelligent-metasurface (SIM)-aided CF [21], and algorithm-unrolling-based distributed optimization for RIS-assisted CF networks [22]; the L4/L3/L2 architectures benchmarked in our testbed directly serve as the wideband physical-layer substrate on which such algorithmic CPU-side processing would be deployed. The consistent numerical finding across this literature is that CF outperforms comparable cMIMO due to favorable-propagation and macro-diversity gains; the present paper is the first to validate the magnitude of these gains experimentally on a wideband, bidirectional, end-to-end testbed.

C. EXPERIMENTAL STUDIES

The work in [23] deployed 32 single-antenna RRUs over a 46.5 m range and performed numerical analysis on the measured channels (channel hardening, aging, SINR), but did not implement a real-time data-transmitting dMIMO system or address fronthaul synchronization. A CF system using commercially available 5G NR RRUs with OTA reciprocity calibration and over-the-air synchronization was demonstrated in [24], but a direct dMIMO-versus-cMIMO comparison was not provided and the commercial hardware does not afford the synchronization fidelity of a custom in-house RoF chain. Centralized and distributed MIMO were experimentally compared at 3.5 GHz indoors in [25], again via numerical post-processing of measured channels rather than active real-time transmission and without an end-to-end spatial-processing framework. Most recently, a large-scale dMIMO platform with real-time adaptive modulation was reported in [26], focused on link-level modulation adaptation for individual UEs without a controlled cMIMO comparison or quantified favorable-propagation/macro-diversity analysis. A detailed comparison with the present work is given in Table 1.

D. REMAINING GAPS AND MAIN CONTRIBUTIONS

To the best of our knowledge, a wideband, multi-AP, multi-antenna, multi-user dMIMO testbed with active bidirectional UL/DL transmission, centralized and distributed spatial-processing algorithms tied to UL-DL duality, a complete end-to-end physical-layer design (packet structure, pilots, wideband OFDM), and a fair direct comparison against cMIMO has not yet been reported in the literature.

Our main contributions, none of which simultaneously hold in any previously reported dMIMO/CF testbed (cf. Table 1), are:

- 1) Wideband (≈ 51 MHz), multi-AP, multi-antenna, multi-user dMIMO operation in both UL and DL

TABLE 1. Comparison of related experimental dMIMO platforms.

Work	Year	Wideband Modeling	Communication Direction		Processing			Experimental Analysis			
			UL	DL	In-house Dev.	Centr. Proc.	Distr. Proc.	d-vs-c Analysis	UL/DL Duality	Sync.	Fav./ Macro.
[25]	2021	×	✓	×	×	×	✓	✓	×	×	×
[23]	2022	×	✓	×	×	×	×	×	×	×	×
[24]	2023	×	✓	✓	×	✓	✓	×	✓	✓	×
[27]	2025	×	✓	×	×	×	✓	×	×	✓	×
[26]	2026	×	×	✓	×	✓	×	×	×	✓	×
This Work	2026	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

Notes: UL: Uplink; DL: Downlink; In-house Dev.: Custom Testbed Development; Centr./Distr. Proc.: Centralized/Distributed Spatial Processing; d-vs-c: Experimental comparison with Collocated MIMO; Sync.: Hardware/Phase Synchronization; Fav./Macro.: Favorable Propagation and Macro-diversity Analysis.

on a single hardware platform, with a complete end-to-end physical-layer chain (preamble, pilots, equalization symbol, payload), as opposed to narrow-band [23], [25] or single-direction [24], [26], [27] setups.

- 2) An in-house dMIMO multi-user, multi-stream testbed built from off-the-shelf components, with SDRoF over optical fiber as fronthaul, achieving sub-10 ps synchronization jitter and 0.45 ps per-hop jitter accumulation, combining the synchronization advantages of analog RoF with the cost efficiency of digital RoF, neither of which is present in commercial-RRU-based [24] nor measured-channel-post-processing [23], [25] studies.
- 3) All three CF processing levels (L4 centralized, L3 hybrid, L2 fully distributed) integrated with UL–DL duality and benchmarked against a cMIMO reference on the *same* hardware, isolating spatial-distribution gains from hardware differences, a direct comparison absent from [23], [24], [25], and [26].
- 4) Bidirectional UL–DL operation with experimentally quantified hardware reciprocity mismatch (10–11 dB), not previously reported in a wideband indoor dMIMO platform.
- 5) Detailed UL and DL timing/frequency synchronization and packet demultiplexing methods for the dMIMO testbed.
- 6) Experimental demonstration of favorable-propagation and macro-diversity gains of dMIMO over cMIMO in both UL and DL, with identification of the spatial-processing methods required for dMIMO to outperform cMIMO in practice.
- 7) Up to 4096-QAM transmission for $K = 1$ and 256-QAM for $K = 4$ in the high-SNR regime on the proposed setup.

The remainder of this paper is organized as follows. Section II describes the system model and the dMIMO architecture. Section III presents the wideband orthogonal frequency division multiplexing (OFDM) signal model for end-to-end UL and DL transmission and discusses three spatial processing algorithms. Section IV presents the sigma-

delta modulation-based dMIMO testbed. Sections V and VI detail the experimental setup and provide an extensive analysis of the experimental results. Section VII provides further discussion on practical implications, limitations, and future directions, and Section VIII concludes the paper.

Notations: Scalars are denoted by italic letters, while vectors and matrices are represented by bold-face lower- and upper-case letters, respectively. For a complex-valued vector $\mathbf{v}$, $\mathbf{v}^T$ denotes the transpose, $\mathbf{v}^*$ denotes the complex conjugate, and $\mathbf{v}^H$ denotes the conjugate transpose. The n -th element of $\mathbf{v}$ is denoted by v_n . The Euclidean norm is represented by $\|\mathbf{v}\|$, and $\text{diag}(\mathbf{v})$ denotes a diagonal matrix with the elements of $\mathbf{v}$ along the main diagonal. $\mathbf{1}_N$ represents a vector of size N with all entries equal to 1. For a complex-valued matrix $\mathbf{M}$, $\mathbf{M}_{(n,m)}$ denotes the entry in the n -th row and m -th column, $\|\mathbf{M}\|_F$ denotes the Frobenius norm, and $\text{tr}(\mathbf{M})$ denotes the trace. The notation $\mathbf{M} \geq 0$ represents a positive semi-definite matrix. The Hadamard (element-wise) product between two matrices is denoted by $\mathbf{A} \odot \mathbf{B}$. For a square matrix $\mathbf{A}$, $\mathbf{I}_A$ and $\mathbf{A}^{-1}$ denote the identity matrix and the inverse matrix, respectively. The operator $\mathbb{E}[\cdot]$ represents statistical expectation, and $\mathcal{CN}(\mu, \sigma^2)$ denotes a circularly symmetric complex Gaussian distribution with mean μ and variance σ^2 .

II. SYSTEM MODEL AND DMIMO ARCHITECTURE

We consider a generalized CF mMIMO wireless communication system, as illustrated in Fig. 1. The network comprises U peripheral CPUs, indexed by $u = 1, \dots, U$, where the symbol U (rather than C) is used for the CPU count to avoid any clash with the subcarrier count C introduced in Section III. Each CPU manages F fronthauls, and each fronthaul connects to A APs. The a -th AP attached to the f -th fronthaul of the u -th CPU is denoted $\text{AP}_{a,f,u}$ and is equipped with $N_{t,a,f,u}$ transmit antennas.

On the user side, there are K UEs. Each UE k is equipped with $N_{r,k}$ receive antennas and supports L_k data streams, for $k = 1, \dots, K$. Consequently, the system-wide totals for transmit antennas, receive antennas, and data streams are defined as $N_t = \sum_{u=1}^U \sum_{f=1}^F \sum_{a=1}^A N_{t,a,f,u}$, $N_r = \sum_{k=1}^K N_{r,k}$, and $L = \sum_{k=1}^K L_k$, respectively. Whenever the CPU/fronthaul




FIGURE 1. System model for CF mMIMO architecture with distributed CPU processing. In each fronthaul, APs are indexed by $a \in \{1, \dots, A\}$; the last AP on every fronthaul bus therefore corresponds to $AP_{A,f,u}$ (the index A rather than the running index a should be read at the terminal AP of every chain). The CPU index u runs over $\{1, \dots, U\}$, and the index c is reserved for the OFDM subcarrier (Section III).




FIGURE 2. Detailed architecture of the AP for UL and DL wideband OFDM processing.

context is clear, we abbreviate $N_{t,a,f,u}$ as $N_{t,a}$ for notational compactness; the full triple-indexed form is invoked only where the multi-CPU/multi-fronthaul dependence is explicitly required. Throughout the rest of the paper, the index c is reserved exclusively for the OFDM subcarrier index introduced in Section III.

In this generalized architecture, processing power is distributed across the network through multiple CPUs, enabling scalability of both fronthaul bandwidth and computation as the number of APs grows.

In Fig. 2, we depict the block diagram of an AP (blue boxes in Fig. 1). Each AP comprises the essential baseband and RF processing blocks for wideband OFDM operation, including channel encoding/decoding, symbol modulation/demodulation (e.g., M -QAM/PSK), spatial pro-

cessing for transmit beamforming, power allocation, and receive combining, OFDM modulation/demodulation via IFFT/FFT, and pulse shaping with an analog front end (AFE) for digital-to-RF interfacing through data conversion, up/down-conversion, and amplification.

In a CF mMIMO wireless network, hardware functional blocks can be partitioned between the CPU and the AP, with significant trade-offs in CPU–AP hardware complexity, signal-processing overhead, and fronthaul bandwidth [5]. Fig. 3 shows the AP-side splits enabling centralized L4 processing: blocks to the left of a split are relocated to the CPU, those to the right stay at the AP.

- *C1 Split:* The spatial processing blocks (e.g., beamforming design), encoding/decoding, and mapping/demapping functions are moved to the CPU, while the remaining physical-layer components are processed at the AP.
- *C2 and C3 Splits:* These configurations progressively shift more functions to the CPU, reducing the AP to a remote radio unit (RRU). In the extreme case of the *C3 split*, the AP consists of simple antennas connected to the CPU via an analog fronthaul.

Regardless of the specific partitioning, the C1, C2, and C3 splits all enable L4 centralized processing. A more detailed analysis of the trade-offs between fronthaul bandwidth and hardware complexity associated with these splits can be found in [5].

Fig. 4 shows the AP-side splits for distributed processing: in L3 (hybrid centralized–distributed), beamforming and combining are computed locally at each AP using local CSI




FIGURE 3. AP functional splits for centralized L4 processing.




FIGURE 4. AP functional splits for L3 and L2 distributed processing.

while the CPU optimizes only the power allocation; in L2 (fully distributed), both the precoders and the power allocation are computed locally at each AP. The corresponding C4 split places the spatial-processing block on the AP side; a comprehensive performance/complexity analysis appears in [5].

The UE architecture for both UL and DL transmission mirrors the AP block diagram shown in Fig. 2, with the processing directions of the UL and DL chains reversed to reflect the user-side perspective.

III. END-TO-END WIDEBAND OFDM SIGNAL MODEL

This section presents the end-to-end signal model for multi-AP, multi-user, multi-stream wideband OFDM transmission under L4, L3, and L2 architectures in both UL and DL.

A. UPLINK

The UL data symbol per subcarrier for the data portion is written (3), as shown at the bottom of the next page, with the time-domain transmit packet (preamble, pilots, equalization symbol, payload) derived in the subsequent subsections. In (3), the trailing subscript $(\cdot)_c$ on each bracketed matrix and vector indicates that all enclosed entries are the per-subcarrier samples evaluated at subcarrier index c . Equivalently, $[\mathbf{f}_{k,1} \ \cdots \ \mathbf{f}_{k,L_k}]_c \triangleq [\mathbf{f}_{k,1,c} \ \cdots \ \mathbf{f}_{k,L_k,c}] \in \mathbb{C}^{N_{r,k} \times L_k}$,

$[\sqrt{q_{k,l,c}}]_c \triangleq \sqrt{q_{k,l,c}}$, and $[s_{k,l}]_c \triangleq s_{k,l,c}$. We adopt this compact notation purely for typographical brevity in displayed matrix expressions; whenever the per-subcarrier dependence is later required without ambiguity, the explicit form $\mathbf{f}_{k,l,c}$ is used. In (3), $s_{k,l,c}$, $q_{k,l,c}$, and $\mathbf{f}_{k,l,c} \in \mathbb{C}^{N_{r,k} \times 1}$ are the frequency-domain data symbol, the allocated transmit power, and the digital UL transmit spatial filter assigned to the l -th data stream of UE k on subcarrier c , respectively. Here, $f_{k,l,c}(n_{r,k})$ denotes the $n_{r,k}$ -th entry of the precoder vector $\mathbf{f}_{k,l,c}$, and we adopt the unit-norm convention $\|\mathbf{f}_{k,l,c}\|_2^2 = 1$ throughout. The power allocation is subject to the following constraint:

$$\sum_{l=1}^{L_k} q_{k,l,c} |f_{k,l,c}(n_{r,k})|^2 \leq q_{k,P_{\max},c,n_{r,k}}, \quad (1)$$

where $q_{k,P_{\max},c,n_{r,k}}$ denotes the maximum transmit power budget allocated to the c -th subcarrier and the $n_{r,k}$ -th antenna of the k -th UE. The use of $|f_{k,l,c}(n_{r,k})|^2$ on the left-hand side correctly captures the per-antenna transmit power because the squared magnitude of the precoder coefficient quantifies the fraction of stream power radiated through the $n_{r,k}$ -th antenna.

In the UL, the distinction between L4, L3, and L2 processing is determined at the APs and the CPU through the implementation of the UL combining method. For example, as discussed before, if the raw time-domain data received at all APs is aggregated at the CPU for global channel estimation and combiner optimization, the system operates under L4 processing, i.e., fully centralized. Conversely, if channel estimation and combiner optimization are performed locally at each AP based on local CSI, the system follows a distributed processing architecture.

Specifically, if these locally estimated data symbols are forwarded to the CPU and simply averaged to produce the final data symbol, it is categorized as L2 distributed processing. If the local estimates are instead combined at the CPU using optimized weights to maximize the combining gain, it is referred to as L3 distributed processing.

By aggregating the spatially projected data symbols across all $N_{r,k}$ antennas and C_a active subcarriers, the frequency-domain representation for the k -th UE at OFDM symbol index m_d can be expressed in matrix form (4), as shown at the bottom of the next page. In (4), the $n_{r,k}$ -th row vector, $\tilde{\mathbf{s}}_{k,UL,m_d}^T = [\tilde{s}_{k,UL,1}(n_{r,k}), \dots, \tilde{s}_{k,UL,C_a}(n_{r,k})]$, represents the frequency-domain signal projected onto the $n_{r,k}$ -th antenna across the active subcarrier range. This matrix $\tilde{\mathbf{S}}_{k,UL,m_d}$ thus captures the complete spatial-frequency state of the k -th user's data portion within a single OFDM symbol.

Following the spatial processing described in (3) and (4), the UE applies IFFT to convert the frequency-domain data symbols into time-domain samples. For the $n_{r,k}$ -th antenna, the l -th time-domain sample of the m_d -th OFDM symbol, we obtain (5), as shown at the bottom of the next page, where C denotes the total number of subcarriers (and corresponding time-domain samples), and $[m_d, l]$ identifies the l -th sample within the m_d -th data OFDM symbol. This transformation




FIGURE 5. UL packet structure.




FIGURE 6. DL packet structure.

provides the time-domain representation of the spatially projected data for UE k at the $n_{r,k}$ -th antenna.

To mitigate inter-symbol interference, a cyclic prefix (CP) of length C_{CP} is added to the time-domain vector. The resulting signal for the m_d -th OFDM symbol, inclusive of the CP, we express as:

$$\bar{\mathbf{s}}_{t,k,UL}^{CP}(n_{r,k})[m_d] = \begin{bmatrix} \bar{\mathbf{s}}_{t,k,UL}(n_{r,k})[m_d, C - C_{CP}] \\ \vdots \\ \bar{\mathbf{s}}_{t,k,UL}(n_{r,k})[m_d, C - 1] \\ \bar{\mathbf{s}}_{t,k,UL}(n_{r,k})[m_d, 0] \\ \vdots \\ \bar{\mathbf{s}}_{t,k,UL}(n_{r,k})[m_d, C - 1] \end{bmatrix}. \quad (2)$$

The total number of samples per OFDM symbol is now $C + C_{CP}$.

By concatenating multiple CP-OFDM data symbols, the discrete-time data portion of the transmitted frame is constructed (6), as shown at the bottom of the page, where




FIGURE 7. One OFDM symbol.

n represents the global discrete-time sample index. While (6) defines the data payload, robust UL communication requires additional reference resources, including a preamble for timing synchronization, pilots for channel estimation, and dedicated sequences for equalization factor estimation.

The comprehensive UL packet structure for a single UE is illustrated in Fig. 5, where Fig. 7 depicts one OFDM symbol

$$\bar{\mathbf{s}}_{k,UL,c} = [\mathbf{f}_{k,1} \cdots \mathbf{f}_{k,L_k}]_c \begin{bmatrix} \sqrt{q_{k,1}} & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & \sqrt{q_{k,L_k}} \end{bmatrix}_c \begin{bmatrix} s_{k,1} \\ \vdots \\ s_{k,L_k} \end{bmatrix}_c \quad (3)$$

$$\bar{\mathbf{s}}_{k,UL,m_d} = \begin{bmatrix} \bar{\mathbf{s}}_{k,UL,m_d}^1 \\ \bar{\mathbf{s}}_{k,UL,m_d}^2 \\ \vdots \\ \bar{\mathbf{s}}_{k,UL,m_d}^{N_{r,k}} \end{bmatrix} = \begin{bmatrix} \bar{s}_{k,UL,1}(1) & \bar{s}_{k,UL,2}(1) & \cdots & \bar{s}_{k,UL,C_a}(1) \\ \bar{s}_{k,UL,1}(2) & \bar{s}_{k,UL,2}(2) & \cdots & \bar{s}_{k,UL,C_a}(2) \\ \vdots & \vdots & \ddots & \vdots \\ \bar{s}_{k,UL,1}(N_{r,k}) & \bar{s}_{k,UL,2}(N_{r,k}) & \cdots & \bar{s}_{k,UL,C_a}(N_{r,k}) \end{bmatrix} \quad (4)$$

$$\bar{\mathbf{s}}_{t,k,UL}(n_{r,k})[m_d, l] = \frac{1}{\sqrt{C}} \sum_{c=0}^{C-1} \bar{\mathbf{s}}_{k,UL,c}(n_{r,k}) e^{j \frac{2\pi c l}{C}}, \quad l = 0, \dots, C - 1 \quad (5)$$

$$\bar{\mathbf{s}}_{t,k,UL}^{tx}(n_{r,k})[n] = \sum_{m_d=1}^{M_d} \bar{\mathbf{s}}_{t,k,UL}^{CP}(n_{r,k})[m_d, n - (m_d - 1)(C + C_{CP})] \quad (6)$$

with active subcarriers shown in red and inactive subcarriers in blue. The frame consists of S total OFDM symbols, where M_d symbols are allocated to the data payload. Furthermore, $K \times N_{r,k}$ symbols serve as channel estimation pilots, while one symbol (typically at the $KN_{r,k} + 1$ index) is reserved for equalization factors. These factors capture the end-to-end effective channel, as detailed in the subsequent sections. Finally, an initial OFDM symbol is dedicated to the preamble to facilitate timing synchronization at the UL receiver.

To generate these reference signals, we utilize Zadoff-chu (ZC) sequences with distinct prime root coefficients. ZC sequences provide constant amplitude zero cyclic auto-correlation (CAZAC) properties, excellent cross-correlation characteristics, and a low peak-to-average power ratio (PAPR) [28]. Specifically, we use three distinct ZC sequences for the preamble, the pilots, and the equalization symbols. Although the frame allocates multiple OFDM symbols for pilots, they remain orthogonal in the time domain as they are dedicated per UE antenna. Consequently, a single base sequence suffices for all pilot symbols within the frame.

B. FREQUENCY-DOMAIN ZC SEQUENCES FOR OFDM SIGNALING

Let C_a denote the number of active subcarriers, where C_a is chosen as a prime number. A ZC sequence of length C_a and root index r , where $\gcd(r, C_a) = 1$, is defined in the frequency domain as:

$$o_r[n] = \exp\left(-j\frac{\pi r n(n+1)}{C_a}\right), \quad n = 0, 1, \dots, C_a - 1. \quad (7)$$

To support the distinct processing tasks within the UL frame, the system generates three ZC sequences using unique prime root indices, $r \in \{r_1, r_2, r_3\}$. Specifically, we assign these roots to the following functional roles:

- **Preamble:** $o_{r_1}[n]$ facilitates timing and frequency synchronization.
- **Pilots:** $o_{r_2}[n]$ enables channel estimation across the $K \times N_{r,k}$ dedicated symbols.
- **Equalization:** $o_{r_3}[n]$ assists in the estimation of end-to-end equalization factors.

The resulting frequency-domain symbol for the i -th reference signal is constructed by mapping the ZC sequence onto the active subcarrier set C_a . The c -th subcarrier of the i -th reference symbol, we express as:

$$X_i[c] = \begin{cases} \sqrt{q_i} o_{r_i}[n], & c \in C_a, \quad n = 0, \dots, C_a - 1 \\ 0, & c \notin C_a \end{cases}, \quad (8)$$

where q_i represents the transmit power allocated to the sequence and c is the subcarrier index.

Following this mapping, an IFFT is applied to generate the time-domain OFDM symbol. Due to the CAZAC property, the frequency-domain ZC sequence preserves its constant amplitude after the IFFT, effectively maintaining a low PAPR for the transmitted reference signals.

C. TIME-DOMAIN MIMO CHANNEL CONVOLUTION

Once the UL packets are prepared for transmission, all K UEs transmit their respective signals simultaneously. These signals undergo a discrete-time convolution with the physical multipath MIMO channel. The signal received at the n_t -th antenna of AP a at discrete-time index n , we express (12), as shown at the bottom of the next page, where $h_{n_t,a,n_r,k}[\tau]$ represents the τ -th tap of the time-domain channel between the n_r -th antenna of UE k and the n_t -th antenna of AP a , with L_h being the total number of channel taps. The term $z_{n_t,a}[n]$ denotes additive white Gaussian noise (AWGN) following a complex normal distribution, $z_{n_t,a}[n] \sim \mathcal{CN}(0, \sigma^2)$.

The signals received across all $N_{t,a}$ antennas of AP a can be organized into a vector as follows:

$$\bar{\mathbf{y}}_{t,a,UL}^{rx}[n] = \begin{bmatrix} \bar{\mathbf{y}}_{t,1,a,UL}^{rx}[n] \\ \vdots \\ \bar{\mathbf{y}}_{t,N_{t,a},a,UL}^{rx}[n] \end{bmatrix}. \quad (9)$$

Furthermore, the collective signal received across all A APs in the network is aggregated into a global receive vector:

$$\bar{\mathbf{y}}_{t,UL}^{rx}[n] = \begin{bmatrix} \bar{\mathbf{y}}_{t,1,UL}^{rx}[n] \\ \vdots \\ \bar{\mathbf{y}}_{t,N_{t,1},1,UL}^{rx}[n] \\ \vdots \\ \bar{\mathbf{y}}_{t,1,A,UL}^{rx}[n] \\ \vdots \\ \bar{\mathbf{y}}_{t,N_{t,A},A,UL}^{rx}[n] \end{bmatrix} = \begin{bmatrix} \left. \begin{matrix} \bar{\mathbf{y}}_{t,1,1,UL}^{rx}[n] \\ \vdots \\ \bar{\mathbf{y}}_{t,N_{t,1},1,UL}^{rx}[n] \end{matrix} \right\} \text{AP 1} \\ \vdots \\ \left. \begin{matrix} \bar{\mathbf{y}}_{t,1,A,UL}^{rx}[n] \\ \vdots \\ \bar{\mathbf{y}}_{t,N_{t,A},A,UL}^{rx}[n] \end{matrix} \right\} \text{AP A} \end{bmatrix}. \quad (10)$$

This vector represents the comprehensive time-domain UL signal received across the entire distributed antenna system at sample n .

D. UL TIMING SYNCHRONIZATION AND FRAME DE-MULTIPLEXING

As discussed before, in the UL, each UE $k \in \{1, \dots, K\}$ transmits a time-domain signal comprising a ZC preamble, followed by orthogonal pilot sequences, an equalization symbol, and the data payload. Due to unknown propagation delays and multipath effects, timing synchronization is performed independently at each antenna of each AP using a correlation-based approach.

1) RECEIVED SIGNAL MODEL FOR SYNCHRONIZATION

The received discrete-time UL signal at the n_t -th antenna of AP a , as established in (12), serves as the input to the synchronization block.

2) CORRELATION-BASED TIMING METRIC

Let $\bar{\mathbf{p}}_t[n]$ denote the known time-domain preamble sequence (root r_1). Timing synchronization is achieved by computing the sliding correlation in (13), as shown at the bottom of the next page, where d represents the discrete-time lag index,

N_p is the preamble length, and N_{tot} is the total number of captured samples.

3) COARSE TIMING OFFSET ESTIMATION

The coarse estimate of the frame start is identified by locating the peak of the correlation metric:

$$\hat{d}_{n_i,a} = \arg \max_d \Lambda_{n_i,a}[d]. \quad (11)$$

The estimated start sample is adjusted by a safety margin δ as $\hat{n}_{start} = \hat{d}_{n_i,a} - \delta$. This ensures the subsequent FFT windows begin within the cyclic prefix, preserving subcarrier orthogonality and avoiding inter-carrier interference.

4) REFERENCE SIGNAL AND PAYLOAD EXTRACTION

Following synchronization, the receiver extracts the different functional components of the frame. Let $S_p = K \times N_{r,k}$ be the number of pilot symbols and $S_e = 1$ be the number of equalization symbols. The symbols are extracted as follows:

a) CHANNEL ESTIMATION PILOTS

For each pilot symbol index $m_p \in \{1, \dots, S_p\}$, the time-domain samples (after CP removal) are (14), as shown at the bottom of the page, where $L_{sym} = C + C_{CP}$ and $n = 0, \dots, C - 1$.

b) EQUALIZATION SYMBOLS

The equalization symbol ($m_e = 1$) is located immediately after the pilot block:

$$\mathbf{y}_{t,n_i,a}^{eq}[n] = \bar{\mathbf{y}}_{t,n_i,a,UL}^{rx}[\hat{n}_{start} + N_p + S_p L_{sym} + C_{CP} + n].$$

c) DATA PAYLOAD

Finally, the M_d data symbols are extracted by jumping over the preamble, pilots, and equalization symbols, in (15), as shown at the bottom of the page, for $m_d = 1, \dots, M_d$. Each extracted symbol is subsequently transformed into the frequency domain via a C -point FFT for further processing.

E. UL MIMO CHANNEL ESTIMATION

Following the successful identification of the frame start sample $\hat{n}_{start}$ via timing synchronization, the receiver extracts the block of $S_p = K \times N_{r,k}$ pilot symbols. For each pilot symbol, the receiver removes the cyclic prefix and performs a C -point FFT to obtain the frequency-domain observation.

Let $Y_{n_i,a,UL}[m_p, c]$ denote the received frequency-domain signal at the n_i -th antenna of AP a on subcarrier c during the

m_p -th pilot OFDM symbol. Since each UE antenna is allocated a unique time-domain symbol index $m_p \in \{1, \dots, S_p\}$, there is no pilot contamination between UEs or their respective antennas. For a given UE k and its $n_{r,k}$ -th antenna, the corresponding pilot symbol index is denoted as $m_p(n_{r,k}, k)$. The least-squares estimate of the channel coefficient, we obtain as:

$$\hat{h}_{n_i,a,UL,n_{r,k},k,c} = \frac{Y_{n_i,a,UL}[m_p, c]}{o_{r_2,k}[c] \sqrt{q_{p,k,c}}}, \quad (16)$$

where $o_{r_2,k}[c]$ is the frequency-domain ZC pilot sequence (root r_2) and $q_{p,k,c}$ represents the transmit power allocated to the pilot subcarrier.

Based on these individual coefficients, the estimated MIMO channel matrix between AP a and UE k on subcarrier c , denoted as $\tilde{\mathbf{H}}_{a,k,c} \in \mathbb{C}^{N_{t,a} \times N_{r,k}}$, is constructed as:

$$\tilde{\mathbf{H}}_{a,k,c} = \begin{bmatrix} \hat{h}_{1,a,UL,1,k,c} & \cdots & \hat{h}_{1,a,UL,N_{r,k},k,c} \\ \vdots & \ddots & \vdots \\ \hat{h}_{N_{t,a},a,UL,1,k,c} & \cdots & \hat{h}_{N_{t,a},a,UL,N_{r,k},k,c} \end{bmatrix}. \quad (17)$$

To facilitate centralized processing across all A APs, the global UL MIMO channel matrix for UE k on subcarrier c is formed by stacking the local channel matrices from all APs:

$$\tilde{\mathbf{H}}_{k,c} = \begin{bmatrix} \tilde{\mathbf{H}}_{1,k,c} \\ \vdots \\ \tilde{\mathbf{H}}_{A,k,c} \end{bmatrix} \in \mathbb{C}^{N_t \times N_{r,k}}, \quad (18)$$

where $N_t = \sum_{a=1}^A N_{t,a}$ is the total number of AP-side antennas across the distributed AP network (which act as receive antennas in the UL). This global channel matrix $\tilde{\mathbf{H}}_{k,c}$ is computed for all UEs $k \in \{1, \dots, K\}$ and all active subcarriers $c \in \mathcal{C}_a$, providing the necessary CSI for the subsequent equalization and detection stages.

F. UL-DL DUALITY-BASED SPATIAL PROCESSING

By utilizing the estimated CSI, i.e., (18), we optimize the UL receive spatial combiners and allocate DL resources for transmission as follows. All three algorithms below exploit the *MIMO SINR uplink-downlink duality* of [29], [30], and [31] and its multi-cell/CF extension in [2] and [3], which states that, under a sum-power constraint, the per-stream SINR region achievable on the broadcast channel coincides with that of its dual multiple-access channel obtained by transposing all channel matrices. Algorithm 1 (L4) applies the global

$$\bar{\mathbf{y}}_{t,n_i,a,UL}^{rx}[n] = \sum_{k=1}^K \sum_{n_r=1}^{N_{r,k}} \sum_{\tau=0}^{L_h-1} h_{n_i,a,n_r,k}[\tau] \bar{\mathbf{s}}_{t,k,UL}^{tx}(n_{r,k})[n - \tau] + z_{n_i,a}[n] \quad (12)$$

$$\Lambda_{n_i,a}[d] = \frac{1}{N_p} \left| \sum_{n=0}^{N_p-1} \bar{\mathbf{y}}_{t,n_i,a,UL}^{rx}[d+n] \bar{\mathbf{p}}_t^*[n] \right|, \quad d = 0, \dots, N_{tot} - N_p \quad (13)$$

$$\mathbf{y}_{t,n_i,a}^{pilot}[m_p, n] = \bar{\mathbf{y}}_{t,n_i,a,UL}^{rx}[\hat{n}_{start} + N_p + (m_p - 1)L_{sym} + C_{CP} + n] \quad (14)$$

$$\mathbf{y}_{t,n_i,a}^{data}[m_d, n] = \bar{\mathbf{y}}_{t,n_i,a,UL}^{rx}[\hat{n}_{start} + N_p + (S_p + S_e + m_d - 1)L_{sym} + C_{CP} + n] \quad (15)$$

Algorithm 1 UL–DL Duality Based L4 Centralized Spatial Processing

- 1: **Initialize:** Set $\kappa = 0$. For all $k \in \{1, \dots, K\}$, $l \in \{1, \dots, L_k\}$, and $c \in \mathcal{C}_a$:
- 2: • $\tilde{\mathbf{H}}_{k,c}$: Global estimated channel; $\mathbf{f}_{k,l,c}^{(0)} \in \mathbb{C}^{N_{r,k} \times 1}$: Initial DL combiner.
- 3: • $q_{k,l,c}^{(0)}$: Initial UL power; $P_{\max}$: Max DL power; ε : Tolerance; $\kappa_{\max}$: Max iterations.
- 4: • σ^2 : Noise variance.[†]
- 5: **repeat**
- 6: **UL Receive Combiner Optimization** ($\mathbf{w}_{k,l,c}$):
- 7: Compute the L4 global receive combiner by solving the generalized eigenvalue problem:

$$\mathbf{w}_{k,l,c}^{(\kappa+1)} = \frac{\lambda_{\max}\{\mathbf{S}_{k,l,c}^{\mathbf{f}}, \mathbf{T}_{k,l,c}^{\mathbf{f}}\}}{\|\lambda_{\max}\{\mathbf{S}_{k,l,c}^{\mathbf{f}}, \mathbf{T}_{k,l,c}^{\mathbf{f}}\}\|} \quad (19)$$

$$\text{where } \mathbf{S}_{k,l,c}^{\mathbf{f}} = q_{k,l,c}^{(\kappa)} \tilde{\mathbf{H}}_{k,c} \mathbf{f}_{k,l,c}^{(\kappa)} (\mathbf{f}_{k,l,c}^{(\kappa)})^H \tilde{\mathbf{H}}_{k,c}^H,$$

$$\begin{aligned} \mathbf{T}_{k,l,c}^{\mathbf{f}} &= \sum_{i \neq k} \sum_{j=1}^{L_i} q_{i,j,c}^{(\kappa)} \tilde{\mathbf{H}}_{i,c} \mathbf{f}_{i,j,c}^{(\kappa)} (\mathbf{f}_{i,j,c}^{(\kappa)})^H \tilde{\mathbf{H}}_{i,c}^H \\ &\quad + \sum_{j \neq l}^{L_k} q_{k,j,c}^{(\kappa)} \tilde{\mathbf{H}}_{k,c} \mathbf{f}_{k,j,c}^{(\kappa)} (\mathbf{f}_{k,j,c}^{(\kappa)})^H \tilde{\mathbf{H}}_{k,c}^H + \sigma^2 \mathbf{I}. \end{aligned} \quad (20)$$

- 8: **DL Receive Combiner Optimization** ($\mathbf{f}_{k,l,c}$):
- 9: Using UL–DL duality, update the DL receive combiners using $\mathbf{w}_{k,l,c}^{(\kappa+1)}$:

$$\mathbf{f}_{k,l,c}^{(\kappa+1)} = \frac{\lambda_{\max}\{\mathbf{S}_{k,l,c}^{\mathbf{w}}, \mathbf{T}_{k,l,c}^{\mathbf{w}}\}}{\|\lambda_{\max}\{\mathbf{S}_{k,l,c}^{\mathbf{w}}, \mathbf{T}_{k,l,c}^{\mathbf{w}}\}\|} \quad (21)$$

$$\text{where } \mathbf{S}_{k,l,c}^{\mathbf{w}} = p_{k,l,c}^{(\kappa)} \tilde{\mathbf{H}}_{k,c}^H \mathbf{w}_{k,l,c}^{(\kappa+1)} (\mathbf{w}_{k,l,c}^{(\kappa+1)})^H \tilde{\mathbf{H}}_{k,c},$$

$$\begin{aligned} \mathbf{T}_{k,l,c}^{\mathbf{w}} &= \sum_{i \neq k} \sum_{j=1}^{L_i} p_{i,j,c}^{(\kappa)} \tilde{\mathbf{H}}_{i,c}^H \mathbf{w}_{i,j,c}^{(\kappa+1)} (\mathbf{w}_{i,j,c}^{(\kappa+1)})^H \tilde{\mathbf{H}}_{i,c} \\ &\quad + \sum_{j \neq l}^{L_k} p_{k,j,c}^{(\kappa)} \tilde{\mathbf{H}}_{k,c}^H \mathbf{w}_{k,j,c}^{(\kappa+1)} (\mathbf{w}_{k,j,c}^{(\kappa+1)})^H \tilde{\mathbf{H}}_{k,c} \\ &\quad + \sigma^2 \mathbf{I}. \end{aligned} \quad (22)$$

- 10: **Power and Resource Allocation:**
- 11: Optimize DL power $p_{k,l,c}^{(\kappa+1)}$ and fairness weights $s_{k,l,c}$ to maximize the objective as per [5].
- 12: **Update Iteration:** $\kappa \leftarrow \kappa + 1$
- 13: **until** Convergence reached or $\kappa > \kappa_{\max}$

[†]In this work, σ^2 is estimated offline by averaging multiple spectrum analyzer measurements taken prior to the experiment with no active transmission, and is assumed constant throughout the measurement campaign.

Algorithm 2 UL–DL Duality Based L3 Distributed Spatial Processing

- 1: **Initialize:** For all $a \in \{1, \dots, A\}$, $k \in \{1, \dots, K\}$, $l \in \{1, \dots, L_k\}$, and $c \in \mathcal{C}_a$:
- 2: • $\tilde{\mathbf{H}}_{a,k,c} \in \mathbb{C}^{N_{t,a} \times N_{r,k}}$: Local channel; $\mathbf{f}_{k,l,c} \in \mathbb{C}^{N_{r,k} \times 1}$: Transmit precoders.
- 3: • $q_{k,l,c}$: Transmit power.
- 4: **Stage I: UL Local Spatial Combining (at AP a)**
- 5: **for** each AP $a = 1$ to A **do**
- 6: Compute the local L3 combiner $\mathbf{w}_{a,k,l,c} \in \mathbb{C}^{N_{t,a} \times 1}$ by solving:

$$\mathbf{w}_{a,k,l,c} = \frac{\lambda_{\max}\{\mathbf{S}_{a,k,l,c}^{\text{local}}, \mathbf{T}_{a,k,l,c}^{\text{local}}\}}{\|\lambda_{\max}\{\mathbf{S}_{a,k,l,c}^{\text{local}}, \mathbf{T}_{a,k,l,c}^{\text{local}}\}\|} \quad (23)$$

$$\text{where } \mathbf{S}_{a,k,l,c}^{\text{local}} = q_{k,l,c} \tilde{\mathbf{H}}_{a,k,c} \mathbf{f}_{k,l,c} \mathbf{f}_{k,l,c}^H \tilde{\mathbf{H}}_{a,k,c}^H,$$

$$\begin{aligned} \mathbf{T}_{a,k,l,c}^{\text{local}} &= \sum_{i \neq k} \sum_{j=1}^{L_i} q_{i,j,c} \tilde{\mathbf{H}}_{a,i,c} \mathbf{f}_{i,j,c} \mathbf{f}_{i,j,c}^H \tilde{\mathbf{H}}_{a,i,c}^H \\ &\quad + \sum_{j \neq l}^{L_k} q_{k,j,c} \tilde{\mathbf{H}}_{a,k,c} \mathbf{f}_{k,j,c} \mathbf{f}_{k,j,c}^H \tilde{\mathbf{H}}_{a,k,c}^H \\ &\quad + \sigma^2 \mathbf{I}_{N_{t,a}}. \end{aligned} \quad (24)$$

- 7: Extract the local soft estimate: $\hat{\mathbf{x}}_{a,k,l,c} = \mathbf{w}_{a,k,l,c}^H \mathbf{Y}_{f,a,c}$.
- 8: Compute effective local gain $\gamma_{a,k,l,c} = \mathbf{w}_{a,k,l,c}^H \tilde{\mathbf{H}}_{a,k,c} \mathbf{f}_{k,l,c} \sqrt{q_{k,l,c}}$ and residual variance $\beta_{a,k,l,c} = \mathbf{w}_{a,k,l,c}^H \mathbf{T}_{a,k,l,c}^{\text{local}} \mathbf{w}_{a,k,l,c}$.
- 9: **end for**
- 10: **Stage II: UL Network-Level Combining (at CPU)**
- 11: Define $\boldsymbol{\gamma}_{k,l,c} = [\gamma_{1,k,l,c}, \dots, \gamma_{A,k,l,c}]^T \in \mathbb{C}^{A \times 1}$ and $\mathbf{B}_{k,l,c} = \text{diag}(\beta_{1,k,l,c}, \dots, \beta_{A,k,l,c})$.
- 12: Compute the CPU weight vector $\mathbf{g}_{k,l,c} \in \mathbb{C}^{A \times 1}$ for aggregation:

$$\mathbf{g}_{k,l,c} = \frac{\lambda_{\max}\{\boldsymbol{\gamma}_{k,l,c} \boldsymbol{\gamma}_{k,l,c}^H, \mathbf{B}_{k,l,c}\}}{\|\lambda_{\max}\{\boldsymbol{\gamma}_{k,l,c} \boldsymbol{\gamma}_{k,l,c}^H, \mathbf{B}_{k,l,c}\}\|} \quad (25)$$

- 13: **Stage III: UL Signal Reconstruction**
- 14: Compute hierarchically combined soft symbol: $\hat{\mathbf{X}}_{k,l,c}^{L3} = \sum_{a=1}^A g_{a,k,l,c}^* \hat{\mathbf{x}}_{a,k,l,c}$.
- 15: Forward $\hat{\mathbf{X}}_{k,l,c}^{L3}$ for equalization, de-mapping, and channel decoding.
- 16: **Stage IV: DL Spatial Power and Resource Allocation**
- 17: Optimize DL power allocations $p_{a,k,l,c}$ based on large-scale fading statistics as described in [2] and [5].

form of this duality at the CPU using the aggregated channel $\tilde{\mathbf{H}}_{k,c}$, computing jointly the UL receive combiners $\mathbf{w}_{k,l,c}$ at the AP side and the UE-side combiners $\mathbf{f}_{k,l,c}$ that, by duality, serve simultaneously as the UL transmit precoders at the UE and as the DL receive combiners at the UE. Algorithms 2 (L3) and

3 (L2) apply the duality *per-AP* on the local channel $\tilde{\mathbf{H}}_{a,k,c}$, treating each AP as an independent receive–transmit pair with its own per-antenna power constraint, in line with the local-precoding interpretation of CF duality in [2, Sec. 5.4]; the resulting *per-AP* combiners are subsequently aggregated at the CPU using either weighted (L3) or arithmetic (L2) combining.

Algorithm 3 UL–DL Duality Based L2 Distributed Spatial Processing

- 1: **Initialize:** For all $a \in \{1, \dots, A\}$, $k \in \{1, \dots, K\}$, $l \in \{1, \dots, L_k\}$, and $c \in \mathcal{C}_a$:
- 2: • $\tilde{\mathbf{H}}_{a,k,c} \in \mathbb{C}^{N_{t,a} \times N_{r,k}}$: Local channel estimates; $\mathbf{f}_{k,l,c} \in \mathbb{C}^{N_{r,k} \times 1}$: Transmit precoders.
- 3: • $q_{k,l,c}$: Transmit power; P_{total} : Total available DL power budget.
- 4: **Stage I: UL Local Spatial Combining (at AP a)**
- 5: **for** each AP $a = 1$ to A **do**
- 6: Compute the local combiner $\mathbf{w}_{a,k,l,c} \in \mathbb{C}^{N_{t,a} \times 1}$ to maximize local SINR:

$$\mathbf{w}_{a,k,l,c} = \frac{\lambda_{\max}\{\mathbf{S}_{a,k,l,c}^{\text{local}}, \mathbf{T}_{a,k,l,c}^{\text{local}}\}}{\|\lambda_{\max}\{\mathbf{S}_{a,k,l,c}^{\text{local}}, \mathbf{T}_{a,k,l,c}^{\text{local}}\}\|} \quad (26)$$

where $\mathbf{S}_{a,k,l,c}^{\text{local}}$ and $\mathbf{T}_{a,k,l,c}^{\text{local}}$ follow the definitions established in Algorithm 2.

- 7: Extract the local soft estimate: $\hat{x}_{a,k,l,c} = \mathbf{w}_{a,k,l,c}^H \mathbf{Y}_{f,a,c}$.
- 8: Forward soft estimates $\hat{x}_{a,k,l,c}$ to the CPU (no additional CSI or variance statistics required).
- 9: **end for**
- 10: **Stage II: UL Network-Level Combining (at CPU)**
- 11: Perform arithmetic averaging of the soft estimates from all A APs, assuming uniform weighting $g_{a,k,l,c} = 1/A$:

$$\hat{X}_{k,l,c}^{L2} = \frac{1}{A} \sum_{a=1}^A \hat{x}_{a,k,l,c} \quad (27)$$

- 12: **Stage III: UL Signal Reconstruction and Processing**
- 13: Forward the averaged symbol $\hat{X}_{k,l,c}^{L2}$ for equalization, demapping, and channel decoding.
- 14: **Stage IV: DL Equal Power Allocation**
- 15: Employ an equal power allocation (EPA) scheme across all users and APs:

$$p_{a,k,l,c} = \frac{P_{\text{total}}}{A \cdot L \cdot |\mathcal{C}_a|}, \quad L \triangleq \sum_{k=1}^K L_k, \quad (28)$$

where $p_{a,k,l,c}$ denotes the power for user k , stream l , on subcarrier c for AP a .

1) CENTRALIZED SPATIAL PROCESSING (L4)

The CPU exploits the global estimated CSI $\tilde{\mathbf{H}}_{k,c}$ to optimize the UL receive combiners $\mathbf{w}_{k,l,c}$ for all UEs, streams, and subcarriers. Additionally, it optimizes the DL receive combiners (the dual of the UL transmit beamforming vectors) and allocates spatial resources, specifically the DL power allocation $p_{k,l,c}$, for L4 transmission. The duality-based centralized spatial processing is detailed in Algorithm 1.

2) HYBRID CENTRALIZED–DISTRIBUTED L3 SPATIAL PROCESSING

Each AP utilizes its local estimated CSI $\tilde{\mathbf{H}}_{a,k,c}$ to optimize the UL receive AP-local combiners $\mathbf{w}_{a,k,l,c}$ for all UEs, streams, and active subcarriers. After generating local soft estimates, each AP sends these to the CPU along with effective channel

statistics. The CPU then performs the final network-level combining. For L3 DL transmission, APs optimize local transmit beamforming vectors, while the CPU performs power allocation $p_{a,k,l,c}$ based on large-scale statistics. This hybrid L3 processing is detailed in Algorithm 2.

3) FULLY DISTRIBUTED L2 SPATIAL PROCESSING

Similar to L3, each AP utilizes its local estimated CSI $\tilde{\mathbf{H}}_{a,k,c}$ to optimize the local UL receive combiners. However, after generating local soft estimates, each AP forwards them to the CPU without additional statistical information. The CPU then performs the final combining via simple arithmetic averaging. For L2 DL transmission, APs optimize local transmit beamforming vectors, but employ equal power allocation in contrast to the L3 scheme. The L2 spatial processing is detailed in Algorithm 3.

G. EQUALIZATION FOR UL L4

Utilizing the optimized receive combiners $\mathbf{w}$, transmit precoders $\mathbf{f}$, and power allocations q , the CPU retrieves the unequalized soft estimates for each subcarrier. For a specific subcarrier c and OFDM symbol index m , the retrieved soft estimate vector $\hat{\mathbf{s}}_{k,c,m} \in \mathbb{C}^{L_k \times 1}$ for UE k is obtained by applying the global combiner $\mathbf{W}_{k,c}$ to the aggregated received signal as follows:

$$\hat{\mathbf{s}}_{k,c,m} = \mathbf{W}_{k,c}^H \sum_{i=1}^K \tilde{\mathbf{H}}_{i,c} \mathbf{F}_{i,c} \mathbf{Q}_{i,c}^{1/2} \mathbf{s}_{i,c,m} + \mathbf{W}_{k,c}^H \mathbf{z}_{c,m}, \quad (29)$$

where:

- $\mathbf{W}_{k,c} = [\mathbf{w}_{k,1,c}, \dots, \mathbf{w}_{k,L_k,c}] \in \mathbb{C}^{N_t \times L_k}$ is the concatenated global L4 combiner matrix for UE k .
- $\tilde{\mathbf{H}}_{i,c} = [\tilde{\mathbf{H}}_{1,i,c}^T, \dots, \tilde{\mathbf{H}}_{A,i,c}^T]^T \in \mathbb{C}^{N_t \times N_{r,i}}$ represents the global estimated CSI for UE i .
- $\mathbf{F}_{i,c} = [\mathbf{f}_{i,1,c}, \dots, \mathbf{f}_{i,L_i,c}] \in \mathbb{C}^{N_{r,i} \times L_i}$ is the UL precoding matrix.
- $\mathbf{Q}_{i,c} = \text{diag}(q_{i,1,c}, \dots, q_{i,L_i,c}) \in \mathbb{R}_+^{L_i \times L_i}$ is the diagonal UL power allocation matrix.
- $\mathbf{z}_{c,m}$ denotes the additive white Gaussian noise (AWGN) vector at the APs.

The resulting vector $\hat{\mathbf{s}}_{k,c,m}$ represents the unequalized data symbols for UE k across all streams $l \in \{1, \dots, L_k\}$ on subcarrier c and OFDM symbol index m .

Next, the CPU estimates the equalization factors for all UEs, streams, and subcarriers using reference pilot symbols transmitted at a dedicated OFDM symbol index ($m = 0$), i.e., the equalization symbol in Fig. 6. As discussed, we employ ZC sequences with a prime root $r = 2$. Let $o_{k,c}$ denote the known ZC value for UE k on subcarrier c . The CPU computes the estimated effective channel gain vector $\mathbf{e}_{k,c} \in \mathbb{C}^{L_k \times 1}$ as:

$$\mathbf{e}_{k,c} = \frac{1}{o_{k,c}} \left(\mathbf{W}_{k,c}^H \sum_{i=1}^K \tilde{\mathbf{H}}_{i,c} \mathbf{F}_{i,c} \mathbf{Q}_{i,c}^{1/2} \mathbf{o}_{i,c} + \mathbf{W}_{k,c}^H \mathbf{z}_{c,0} \right), \quad (30)$$

The resulting vector $\mathbf{e}_{k,c} = [e_{k,1,c}, \dots, e_{k,L_k,c}]^T \in \mathbb{C}^{L_k \times 1}$ captures the estimated composite gains for each stream, where $\mathbf{o}_{i,c} = o_{i,c} \mathbf{1}_{L_i} \in \mathbb{C}^{L_i \times 1}$ is the per-stream ZC pilot

vector for UE i at subcarrier c (i.e., the same scalar ZC value $o_{i,c}$ broadcast on all L_i streams during the equalization OFDM symbol). While $e_{k,l,c}$ is dominated by the desired term $\mathbf{w}_{k,l,c}^H \tilde{\mathbf{H}}_{k,c} \mathbf{f}_{k,l,c} \sqrt{q_{k,l,c}}$, it also inherently includes residual interference and processed noise from the pilot phase, which impacts the subsequent equalization accuracy.

Finally, the CPU equalizes the soft estimates $\hat{\mathbf{s}}_{k,c,m}$. The equalized symbol vector $\tilde{\mathbf{s}}_{k,c,m} \in \mathbb{C}^{L_k \times 1}$ is retrieved as:

$$\tilde{\mathbf{s}}_{k,c,m} = (\text{diag}(\mathbf{e}_{k,c}))^{-1} \hat{\mathbf{s}}_{k,c,m}, \quad (31)$$

where for each stream l , $\tilde{s}_{k,l,c,m} = e_{k,l,c}^{-1} \hat{s}_{k,l,c,m}$. Given the equalized symbols $\tilde{\mathbf{s}}_{k,c,m}$ and transmitted symbols $\mathbf{s}_{k,c,m}$, the SINDR can be derived as follows.

Following the equalization process, the retrieved symbol $\tilde{s}_{k,l,c,m}$ for the l -th stream of UE k on subcarrier c and OFDM symbol m can be expressed as:

$$\tilde{s}_{k,l,c,m} = s_{k,l,c,m} + \epsilon_{k,l,c,m}, \quad (32)$$

where $s_{k,l,c,m}$ denotes the transmitted constellation point and $\epsilon_{k,l,c,m}$ represents the residual error. This error term encompasses the combined effects of multi-user interference, processed thermal noise, and distortions stemming from hardware impairments and imperfect channel estimation. The error is assumed to be uncorrelated with the transmitted signal, satisfying $\mathbb{E}[s_{k,l,c,m} \epsilon_{k,l,c,m}^*] = 0$.

Accordingly, the error vector magnitude (EVM) for the l -th stream is defined as the root-mean-square of the error normalized by the average signal power:

$$\text{EVM}_{k,l,c} = \sqrt{\frac{\mathbb{E}[|\tilde{s}_{k,l,c,m} - s_{k,l,c,m}|^2]}{\mathbb{E}[|s_{k,l,c,m}|^2]}}. \quad (33)$$

Here, we assume the transmitted symbols are drawn from a normalized constellation with unit average power, i.e., $\mathbb{E}[|s_{k,l,c,m}|^2] = 1$, which is enforced by the power allocation $q_{k,l,c}$. Under this normalization, the SINDR simplifies to:

$$\text{SINDR}_{k,l,c} = \frac{\mathbb{E}[|s_{k,l,c,m}|^2]}{\mathbb{E}[|\tilde{s}_{k,l,c,m} - s_{k,l,c,m}|^2]} = \frac{1}{\text{EVM}_{k,l,c}^2}, \quad (34)$$

which is expressed in decibels (dB) as follows:

$$\begin{aligned} \text{SINDR}_{\text{dB}} &= 10 \log_{10} (\text{SINDR}_{k,l,c}) \\ &= -20 \log_{10} (\text{EVM}_{k,l,c}). \end{aligned} \quad (35)$$

H. EQUALIZATION FOR UL L3

In L3, UL combining optimization is performed at two levels, as detailed in Algorithm 2. First, each AP performs local estimation of the UEs' data symbols using local channel information; second, these estimates are forwarded to the CPU, where they are aggregated using optimal weights.

Specifically, the local estimates of a UE's data at AP a are calculated (36), as shown at the bottom of the next page, where $\mathbf{W}_{a,k,c} = [\mathbf{w}_{a,k,1,c}, \dots, \mathbf{w}_{a,k,L_k,c}] \in \mathbb{C}^{N_{\text{td}} \times L_k}$ is the local combiner matrix at AP a for UE k on subcarrier c , with column vectors $\mathbf{w}_{a,k,l,c}$ defined as in Algorithm 2. Similarly, for a specific UE k and stream l , the local unequalized

estimates from all A APs are aggregated at the CPU (37), as shown at the bottom of the next page.

Using a similar method, the local equalization factors for AP a are estimated using the ZC sequence $o_{k,c}$ (38), as shown at the bottom of the next page.

Finally, the equalized UL data symbols for L3 are retrieved (39), as shown at the bottom of the next page. Given the equalized symbols $\tilde{\mathbf{s}}_{k,l,c,m}^{L3}$ and the transmitted symbols $\mathbf{s}_{k,l,c,m}$, the SINDR and EVM can be calculated using the same methodology discussed for the L4 UL.

In (40), as shown at the bottom of the next page, $\mathbf{E}_{k,c}^{L3} \in \mathbb{C}^{L_k \times L_k}$ is the diagonal effective-gain matrix whose l -th diagonal entry is $e_{k,l,c}^{L3}$.

I. EQUALIZATION FOR UL L2

The L2 UL processing architecture similarly utilizes a two-stage hierarchical approach. However, in L2, the second stage of combining at the CPU is simplified. Instead of using optimized weights as in L3, the CPU performs a simple arithmetic average of the local estimates received from the APs.

The local estimates $\hat{\mathbf{s}}_{k,c,m}^{\text{local},a}$ are generated at each AP a identically to the L3 process. These estimates are aggregated at the CPU to form the global L2 estimate matrix $\hat{\mathbf{S}}_{k,c,m}^{L2}$ as (41), as shown at the bottom of the page 13, where $\mathbf{1}_{k,l}^{\text{CPU}}$ represents a vector of ones of size $A \times 1$. This structure effectively enforces a uniform averaging of the local estimates for each stream.

Similarly, the local equalization factors $\mathbf{e}_{k,c}^{\text{local},a}$ estimated at each AP are aggregated at the CPU to derive the global L2 equalization matrix $\mathbf{E}_{k,c}^{L2}$ (42), as shown at the bottom of the page 13.

Finally, the equalized UL data symbols for L2 are retrieved (43), as shown at the bottom of the page 13. We calculate SINDR and EVM values for the L2 UL, following the same method used in the L4 and L3 cases.

A visual summary of the whole UL transmission and reception is depicted in Fig. 8.

Here $\mathbf{E}_{k,c}^{L2} \in \mathbb{C}^{L_k \times L_k}$ is the diagonal effective-gain matrix.

J. DOWNLINK

We obtain the DL wideband ofdm signal model from the UL framework via UL-DL duality. Specifically, under TDD reciprocity, the DL channel matrix satisfies $\mathbf{H}_{k,c}^{\text{DL}} \approx \tilde{\mathbf{H}}_{k,c}^T \in \mathbb{C}^{N_{r,k} \times N_{t_i}}$. In practice, however, hardware impairments introduce a mismatch, and we model the DL channel as:

$$\mathbf{H}_{k,c}^{\text{DL}} = \mathbf{D}_{\text{RX},k} \tilde{\mathbf{H}}_{k,c}^T \mathbf{D}_{\text{TX}} + \Delta \mathbf{H}_{k,c}, \quad (44)$$

where $\mathbf{D}_{\text{RX},k} \in \mathbb{C}^{N_{r,k} \times N_{r,k}}$ captures the frequency-dependent RX-chain mismatch on the UE- k side, $\mathbf{D}_{\text{TX}} \in \mathbb{C}^{N_{t_i} \times N_{t_i}}$ captures the AP-side TX-chain mismatch, both diagonal, $\Delta \mathbf{H}_{k,c}$ denotes the residual mismatch error, and we recover perfect reciprocity when $\mathbf{D}_{\text{RX},k} = \mathbf{D}_{\text{TX}} = \mathbf{I}$. We identify this mismatch as the primary source of the UL-DL SINDR degradation

quantified experimentally in Section VI. The complete end-to-end DL signal models for L4, L3, and L2 processing levels, including spatial projection, IFFT, CP addition, UE reception, equalization, and retrieval, are provided in the supplementary material.

IV. SIGMA-DELTA MODULATION BASED DMIMO TESTBED

A. SIGMA-DELTA MODULATOR

A sigma-delta modulator (SDM) is a control loop built around a typically very coarse quantizer (often only 1-bit), running at a sampling frequency much higher than the Nyquist rate of the input signal [32]. A block diagram of a single-bit first-order SDM is shown in Fig. 9, where an N -bit input signal $u[n]$ is quantized to a single-bit output signal $v[n]$. This operation will introduce significant quantization noise, which can be modeled as an additive noise

source $e[n]$ at the output of the quantizer. The difference between the input of the quantizer ($s[n]$) and the output of the quantizer ($v[n]$)

$$\begin{aligned} w[n] &= s[n] - v[n] \\ &= s[n] - (s[n] + e[n]) \\ &= -e[n] \end{aligned} \quad (45)$$

is shaped by the loop filter such that most of the quantization noise power is pushed out-of-band. Using the linearized model of the quantizer as an additive noise source, the output signal $V(z^{-1})$ can be expressed in function of the input signal $U(z^{-1})$ and the quantization noise $E(z^{-1})$ (which represent the Z-transforms) as

$$\begin{aligned} V(z^{-1}) &= S(z^{-1}) + E(z^{-1}) \\ &= U(z^{-1}) + z^{-1}W(z^{-1}) + E(z^{-1}) \\ &= U(z^{-1}) - z^{-1}E(z^{-1}) + E(z^{-1}) \end{aligned}$$

$$\hat{\mathbf{s}}_{k,c,m}^{local,a} = \begin{bmatrix} \hat{s}_{k,1,c,m}^{local,a} \\ \hat{s}_{k,2,c,m}^{local,a} \\ \vdots \\ \hat{s}_{k,l,c,m}^{local,a} \\ \vdots \\ \hat{s}_{k,L_k,c,m}^{local,a} \end{bmatrix} = \mathbf{W}_{a,k,c}^H \sum_{i=1}^K \tilde{\mathbf{H}}_{i,a,c} \mathbf{F}_{i,c} \mathbf{Q}_{i,c}^{1/2} \mathbf{s}_{i,c,m} + \mathbf{W}_{a,k,c}^H \mathbf{z}_{a,c,m}, \quad (36)$$

$$\hat{\mathbf{s}}_{k,c,m}^{L3} = \begin{bmatrix} \hat{s}_{k,1,c,m}^{L3} \\ \hat{s}_{k,2,c,m}^{L3} \\ \vdots \\ \hat{s}_{k,l,c,m}^{L3} \\ \vdots \\ \hat{s}_{k,L_k,c,m}^{L3} \end{bmatrix} = \begin{bmatrix} \hat{s}_{k,1,c,m}^{local,1} & \hat{s}_{k,1,c,m}^{local,2} & \dots & \hat{s}_{k,1,c,m}^{local,A} \\ \hat{s}_{k,2,c,m}^{local,1} & \hat{s}_{k,2,c,m}^{local,2} & \dots & \hat{s}_{k,2,c,m}^{local,A} \\ \vdots & \vdots & \ddots & \vdots \\ \hat{s}_{k,l,c,m}^{local,1} & \hat{s}_{k,l,c,m}^{local,2} & \dots & \hat{s}_{k,l,c,m}^{local,A} \\ \vdots & \vdots & \ddots & \vdots \\ \hat{s}_{k,L_k,c,m}^{local,1} & \hat{s}_{k,L_k,c,m}^{local,2} & \dots & \hat{s}_{k,L_k,c,m}^{local,A} \end{bmatrix} [\mathbf{w}_{k,1}^{CPU} \quad \mathbf{w}_{k,2}^{CPU} \quad \dots \quad \mathbf{w}_{k,L_k}^{CPU}] \quad (37)$$

$$\mathbf{e}_{k,c}^{local,a} = \begin{bmatrix} e_{k,1,c}^{local,a} \\ e_{k,2,c}^{local,a} \\ \vdots \\ e_{k,l,c}^{local,a} \\ \vdots \\ e_{k,L_k,c}^{local,a} \end{bmatrix} = \frac{1}{O_{k,c}} \left[\mathbf{W}_{a,k,c}^H \sum_{i=1}^K \tilde{\mathbf{H}}_{i,a,c} \mathbf{F}_{i,c} \mathbf{Q}_{i,c}^{1/2} \mathbf{o}_{i,c} + \mathbf{W}_{a,k,c}^H \mathbf{z}_{a,c,0} \right] \quad (38)$$

$$\mathbf{E}_{k,c}^{L3} = \begin{bmatrix} \mathbf{e}_{k,1,c}^{L3} \\ \mathbf{e}_{k,2,c}^{L3} \\ \vdots \\ \mathbf{e}_{k,l,c}^{L3} \\ \vdots \\ \mathbf{e}_{k,L_k,c}^{L3} \end{bmatrix} = \begin{bmatrix} e_{k,1,c}^{local,1} & e_{k,1,c}^{local,2} & \dots & e_{k,1,c}^{local,A} \\ e_{k,2,c}^{local,1} & e_{k,2,c}^{local,2} & \dots & e_{k,2,c}^{local,A} \\ \vdots & \vdots & \ddots & \vdots \\ e_{k,l,c}^{local,1} & e_{k,l,c}^{local,2} & \dots & e_{k,l,c}^{local,A} \\ \vdots & \vdots & \ddots & \vdots \\ e_{k,L_k,c}^{local,1} & e_{k,L_k,c}^{local,2} & \dots & e_{k,L_k,c}^{local,A} \end{bmatrix} [\mathbf{w}_{k,1}^{CPU} \quad \mathbf{w}_{k,2}^{CPU} \quad \dots \quad \mathbf{w}_{k,L_k}^{CPU}] \quad (39)$$

$$\hat{\mathbf{s}}_{k,c,m}^{L3} = (\mathbf{E}_{k,c}^{L3})^{-1} \hat{\mathbf{s}}_{k,c,m}^{L3} = \begin{bmatrix} e_{k,1,c}^{L3} & 0 & \dots & 0 \\ 0 & e_{k,2,c}^{L3} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & e_{k,L_k,c}^{L3} \end{bmatrix}^{-1} \begin{bmatrix} \hat{s}_{k,1,c,m}^{L3} \\ \hat{s}_{k,2,c,m}^{L3} \\ \vdots \\ \hat{s}_{k,L_k,c,m}^{L3} \end{bmatrix} \quad (40)$$

$$= U(z^{-1}) \cdot \underbrace{1}_{\text{STF}} + E(z^{-1}) \underbrace{(1 - z^{-1})}_{\text{NTF}} \quad (46)$$

which uses the result from (45) to decompose the output into a signal transfer function (STF) and a noise transfer function (NTF). (46) shows the shaping that occurs on the quantization noise by the NTF, which has a high-pass characteristic. The frequency domain representation of this is shown in Fig. 10a where the contributions of an oversampled input signal ($f_s/2$ much larger than the signal bandwidth) and the shaped quantization noise are visible in the output spectrum. From the figure, it should be clear how increasing the oversampling ratio (OSR) (by increasing f_s) or the order of the loop filter (which increases the slope of the NTF) will reduce the residual in-band quantization noise power.

Fig. 10 shows all the steps of a SDM-based all-digital transmitter [33]. After sigma-delta modulation (Fig. 10a), the resulting 1-bit signal can be up-converted to a carrier frequency f_c , e.g., four times the modulator sampling frequency as shown in Fig. 10b. In this case, the required local oscillators for the in-phase and quadrature signals (which are $\pi/2$ phase shifted with respect to each other) can be represented by the digital sequences $\cos(2\pi k f_c / f_s) = \cos(\pi/2 \cdot k) = \{1, 0, -1, 0\}$ and $\sin(2\pi k f_c / f_s) = \{0, 1, 0, -1\}$ respectively as shown in Fig. 11. The result of this operation remains a single-bit signal which can afterward be filtered by an analog band-pass filter around f_c to obtain the final RF signal as shown in Fig. 10c. Amplification of the up-converted signal can be done using non-linear power amplifiers without

introducing significant distortion before the band-pass filter. This operation forms the basis of sigma-delta modulated signals over fiber (SDoF) transmitters as it allows generating high-frequency RF signals using only digital circuits and a simple analog filter.

SDM-based transmission trades the high linearity requirements of traditional multi-bit digital-to-analog converter (DAC) and power-amplifier architectures for higher sample-rate requirements on the digital side. Where a typical baseband only requires the bandwidth of the signal itself, SDM-based transmitters require the bandwidth of the oversampled SDM output signal, which can be multiple times the signal bandwidth.

The single-bit SDM is specifically chosen, rather than a multi-bit alternative, because a single-bit output is inherently free of DAC integral- and differential-non-linearity errors and removes the need for a high-resolution DAC or precision resistor ladder at the remote radio unit (RRU) (the dominant cost and power source in the equivalent multi-bit digital radio-over-fiber (DRoF) transmitter). The single-bit stream can also drive a non-linear, low-cost power amplifier directly without introducing in-band distortion, since its information is encoded in transition timing rather than amplitude. A multi-bit modulator would relax the oversampling ratio but at the cost of a multi-bit DAC per RRU and a more linear optical link, which inverts the cost-vs-bandwidth trade-off targeted by the in-house dMIMO testbed. The single-bit choice therefore reduces the per-RRU hardware cost, which is the dominant cost since each AP needs its own RRU, in exchange for a higher central-unit sample rate.

$$\hat{\mathbf{s}}_{k,c,m}^{L2} = \begin{bmatrix} \hat{s}_{k,1,c,m}^{L2} \\ \hat{s}_{k,2,c,m}^{L2} \\ \vdots \\ \hat{s}_{k,l,c,m}^{L2} \\ \vdots \\ \hat{s}_{k,L_k,c,m}^{L2} \end{bmatrix} = \begin{bmatrix} \hat{s}_{k,1,c,m}^{local,1} & \hat{s}_{k,1,c,m}^{local,2} & \cdots & \hat{s}_{k,1,c,m}^{local,A} \\ \hat{s}_{k,2,c,m}^{local,1} & \hat{s}_{k,2,c,m}^{local,2} & \cdots & \hat{s}_{k,2,c,m}^{local,A} \\ \vdots & \vdots & \ddots & \vdots \\ \hat{s}_{k,l,c,m}^{local,1} & \hat{s}_{k,l,c,m}^{local,2} & \cdots & \hat{s}_{k,l,c,m}^{local,A} \\ \vdots & \vdots & \ddots & \vdots \\ \hat{s}_{k,L_k,c,m}^{local,1} & \hat{s}_{k,L_k,c,m}^{local,2} & \cdots & \hat{s}_{k,L_k,c,m}^{local,A} \end{bmatrix} \begin{bmatrix} \mathbf{1}_{k,1}^{CPU} & \mathbf{1}_{k,2}^{CPU} & \cdots & \mathbf{1}_{k,L_k}^{CPU} \\ A & A & & A \end{bmatrix} \quad (41)$$

$$\mathbf{E}_{k,c}^{L2} = \begin{bmatrix} e_{k,1,c}^{local,1} & e_{k,1,c}^{local,2} & \cdots & e_{k,1,c}^{local,A} \\ e_{k,2,c}^{local,1} & e_{k,2,c}^{local,2} & \cdots & e_{k,2,c}^{local,A} \\ \vdots & \vdots & \ddots & \vdots \\ e_{k,l,c}^{local,1} & e_{k,l,c}^{local,2} & \cdots & e_{k,l,c}^{local,A} \\ \vdots & \vdots & \ddots & \vdots \\ e_{k,L_k,c}^{local,1} & e_{k,L_k,c}^{local,2} & \cdots & e_{k,L_k,c}^{local,A} \end{bmatrix} \begin{bmatrix} \mathbf{1}_{k,1}^{CPU} & \mathbf{1}_{k,2}^{CPU} & \cdots & \mathbf{1}_{k,L_k}^{CPU} \\ A & A & & A \end{bmatrix} \quad (42)$$

$$\hat{\mathbf{s}}_{k,c,m}^{L2} = (\mathbf{E}_{k,c}^{L2})^{-1} \hat{\mathbf{s}}_{k,c,m}^{L2} = \begin{bmatrix} e_{k,1,c}^{L2} & 0 & \cdots & 0 \\ 0 & e_{k,2,c}^{L2} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & e_{k,L_k,c}^{L2} \end{bmatrix}^{-1} \begin{bmatrix} \hat{s}_{k,1,c,m}^{L2} \\ \hat{s}_{k,2,c,m}^{L2} \\ \vdots \\ \hat{s}_{k,l,c,m}^{L2} \\ \vdots \\ \hat{s}_{k,L_k,c,m}^{L2} \end{bmatrix} \quad (43)$$




FIGURE 8. Summary of the UL transmission and reception signal processing flow.




FIGURE 9. Block diagram of a single-bit first-order sigma-delta modulator using a two-level quantizer. The linearized model of the quantizer is shown as an additive noise source $e[n]$.




FIGURE 10. Frequency domain representation of the signals in a sigma-delta modulator with $f_s = 2f_s = 4f_c$. (a) Baseband input signal and shaped quantization noise processed at f_s at the output of the SDM. (b) Up-mixed output signal to carrier frequency f_c . (c) Filtered output signal around f_c after analog band-pass filtering.

B. RADIO-OVER-FIBER SCHEMES

To understand the benefits of using SDM in a RoF system, it is important to first understand the different RoF schemes [34] as shown in Fig. 12. Traditional macro-cellular wireless communications primarily employ DRoF links [35] for their immunity to nonlinear distortion; however, DRoF systems exhibit low bandwidth efficiency [36] and high implementation complexity, as illustrated in Fig. 12a. In Fig. 1 the architecture consists of a central processing unit (CPU) connected to multiple access points (APs). To reduce

hardware complexity, multiple APs are combined in a single RRU who communicate with the CPU through the fronthaul. DRoF links require up-conversion at the RRU to generate the RF signal, which requires a high-frequency local oscillator which should be synchronized across all RRUs in the case of multiple-input multiple-output (MIMO) operation.

Analog radio-over-fiber (ARoF) as illustrated in Fig. 12b alleviates the need for up-conversion at the RRU by transmitting the RF signal directly over the optical fiber. While this




FIGURE 11. Block diagram of the up-conversion of a baseband sigma-delta modulated signal to a carrier frequency f_c using digital mixing.




FIGURE 12. Examples of different radio-over-fiber schemes.

provides inherent synchronization between multiple RRUs and simple RRU architecture, it suffers from high linearity requirements on the optical link to avoid distortion of the transmitted RF signal.

Leveraging the SDM discussed before, the analog up-converted RF signal can be generated at the CPU as shown in Fig. 12c by modulating the n -bit baseband signal into a single-bit SDM output which is afterward up-converted to the carrier frequency while maintaining the single-bit nature of the signal (Fig. 10). This 1-bit up-converted signal can then be transmitted over the optical fiber to the RRU using the same hardware as used in current DRoF systems. At the RRU, the received optical signal is converted back to the electrical domain and the out-of-band quantization noise is filtered using a band-pass filter around the carrier frequency to obtain the final RF signal. As no oscillator is required at the RRU, synchronization between multiple RRUs is inherently achieved as in ARoF systems, while the high

linearity requirements of ARoF systems are alleviated by the use of a 1-bit signal over the optical link.

C. TESTBED ARCHITECTURE

The full SDoF-based MIMO testbed system architecture is shown in Fig. 13 [37]. The testbed consists of a computer acting as CPU with a high-speed field programmable gate array (FPGA) (HTG-930) as distributed unit (DU) connected via the peripheral component interconnect express (PCIe) interface. The CPU and DU act as a single unit in this architecture. The DU is connected to four RRUs using dedicated fiber links for the up- and DL respectively.

Each RRU can transmit or receive four antenna streams using time division duplex (TDD) operation. To transmit or receive these four streams, the SDM data is interleaved before transmission over the optical fiber. An alignment header is added to the data stream to allow for proper deinterleaving.

1) DOWNLINK

Two RRUs are connected to a single FPGA (VCU108) by means of a mezzanine connector to increase the hardware density and while they share the FPGA fabric and high-speed transceivers, each RRU behaves independently. Each RRU receives the incoming optical SDM signal from the quad small form-factor pluggable (QSFP28) transceiver, which also provides a recovered clock used for synchronization which will be discussed later. Each RRU contains four antenna streams which are received as an interleaved SDM data stream. The incoming data is deinterleaved using the alignment header and each antenna stream is digitally up-converted to the carrier frequency as shown in Fig. 11 before leaving the FPGA fabric. An in-house developed active antenna unit (AAU) mezzanine board filters the out-of-band quantization noise and amplifies the RF signal before transmission over the air. Specifically, the AAU implements an explicit RF band-pass filter centered at the carrier frequency f_c [37]. This filter is therefore an integral part of the SDRoF transmitter rather than an optional add-on, and its cost is offset by the elimination of multi-bit GSps DACs and matched RF reconstruction filters that would otherwise be required in a comparable digital RoF link, as discussed in Section VII-B.

2) UPLINK

Using TDD operation, the same RRU hardware is used for the UL. The AAU mezzanine board now receives the RF signal over the air, amplifies it before down-conversion to baseband using a mixer. The analog-to-digital converter (ADC) samples the baseband signal of each antenna stream after which each stream is sigma-delta modulated as shown before. The resulting SDM streams of the four antenna signals are interleaved using the alignment header before transmission over the optical fiber to the DU.

3) SYNCHRONIZATION

The synchronization circuit implemented on the AAU is discussed in detail in [38] but repeated here for




FIGURE 13. Testbed system architecture with a single DU connected to four RRUs using a pair (up/down link) of dedicated fibers per RRU. As the DU is connected to the CPU using PCIe, the CPU and DU act as a single unit. (QSPF: quad small form-factor pluggable; PCIe: peripheral component interconnect express; FPGA: field-programmable gate array; AAU: active antenna unit; SDM: sigma-delta modulator; PLL: phase-locked loop; ADC: analog-to-digital converter.)




FIGURE 14. Synchronization circuit implemented on the RRU how the ADC sampling clock and down-converter oscillator is derived from the incoming SDM signal. (ADC: analog-to-digital converter; PLL: phase lock loop.)

completeness. As shown in Fig. 14, the recovered clock from the QSFP28 transceiver is sent to the AAU where a phase locked loop (PLL) generates the ADC sampling clock. This PLL is configured with a very narrow loop bandwidth which provides low-frequency phase synchronization w.r.t. the CPU but heavily filters high-frequency jitter introduced by the QSFP28 transceiver and FPGA fabric clock distribution. Any jitter on the ADC sampling clock will directly translate into signal-to-noise ratio (SNR) degradation as

$$\Delta \text{SNR} = -20 \log(\sigma_{\phi, \text{jitter}}) \quad (47)$$

where $\sigma_{\phi, \text{jitter}}$ is the RMS phase noise (in radians) of the sampling clock. The ADC sampling clock is then used as the input reference for a second PLL generating the carrier frequency clock required for down-conversion. A low phase-noise clock source is required for the down-conversion to

ensure good time synchronization between multiple RRUs for MIMO operation. This architecture combined with the use of SDM makes the whole architecture transparent with respect to the CPU.

Fiber-length mismatch between APs does not affect the achievable E2E gains in this architecture: any fixed propagation-delay difference along the fiber appears as a constant per-AP phase offset at the carrier frequency and is absorbed by the per-frame pilot-based channel estimate, while the recovered-clock jitter at the AP-side PLL is set by the QSFP CDR (chromatic dispersion at 1310 nm in OM4 multi-mode fiber is essentially zero), so unequal fiber lengths across APs do not enlarge the dynamic jitter budget either. The synchronization behaviour observed with 100 m of OM4 multi-mode fiber per AP is therefore representative of setups in which the fiber-length variation between APs is held within a comparable range.

This transparency can be exploited to emulate any functional split introduced in Fig. 3 and Fig. 4 as all signals are fully computed in the CPU before being sent to different APs. In the case of L4 processing, all information is utilized at the CPU, while L2/L3 processing can be emulated by preparing the required antenna streams using only the CSI that would be available at the relevant AP.

V. EXPERIMENTAL SETUP

We consider a multi-user, multi-antenna, multi-AP dMIMO experimental wireless network, as shown in the photograph in Fig. 16. A schematic representation is provided in Fig. 15. The architecture consists of a CPU connected via four optical fiber fronthaul links to four separate processing units. Each




FIGURE 15. Schematic of the experimental setup illustrating the CPU, fronthaul links, and RRU configurations.

TABLE 2. Key experimental and system parameters.

Parameter	Symbol	Value / Setting
Carrier Frequency	f_c	3.5 GHz
Sampling Rate	f_s	57.6 MHz
Occupied Bandwidth	B	≈ 51 MHz
No. of Subcarriers	C	128 (113 active)
CP Length	C_{CP}	32 samples (≈ 556 ns)
Modulation Scheme	—	64-QAM
Channel Coding	—	Convolutional (rate 1/2)
Decoding Algorithm	—	Viterbi
Packet Size	—	1500 Bytes
Total TX Antennas	N_t	16 (4 APs $\times$ 4 ant.)
Active APs per Experiment	A	2 (AP ₁ and AP ₄)
Active TX Antennas	$N_{t,act}$	8 (2 APs $\times$ 4 ant.)
No. of UEs	K	Up to 4 (single-antenna)
AP Separation (dMIMO)	d_{AP}	3 m
AP Spacing (cMIMO)	—	$\lambda/2$ ULA
UE Spacing (collocated)	d_k	≈ 15 cm
AP-to-UE Distance	—	≈ 3.5 m
Max. TX Power/Antenna	P_{max}	-3 dBm
Fronthaul	—	$4 \times$ optical (SDRoF)
Spatial Processing	—	L4, L3, L2
Communication Direction	—	UL and DL (TDD)

unit contains an FPGA for local processing and two RRUs, and each RRU is connected to four antennas, resulting in a total of 16 antennas across four APs on the AP side. The UE side consists of one hardware unit equipped with two RRUs and one FPGA, supporting a maximum of four single-antenna UEs or two dual-antenna UEs.

Unless otherwise specified, the experimental parameters are as follows. In the dMIMO case, AP₁ and AP₄ are active and spaced 3 m apart, whereas in the cMIMO case, the same two APs are collocated with $\lambda/2$ antenna spacing to form a ULA. We employ 64-QAM modulation with a coding rate of 1/2 and Viterbi decoding. The system uses a packet length of 1500 Bytes and an OFDM symbol length of 128 subcarriers (113 active). MMSE digital filters are applied at both the TX and RX sides, with max-min fairness as the objective for power allocation. ZC sequences with the first three prime root indices are utilized for the preamble, equalization symbol, and pilot-based channel estimation. The sampling rate is 57.6 MHz at a carrier frequency of 3.5 GHz. The AP-to-UE distance is approximately 3.5 m, while the UEs are spaced approximately 15 cm apart in the collocated scenario. Detailed parameters are provided in Table 2.

To ensure a fair comparison between dMIMO and cMIMO, the transmit power per antenna is kept identical in both configurations. The sole architectural difference is that in dMIMO the APs are geographically separated by up to 3 m, whereas in cMIMO the same APs are collocated with $\lambda/2$ spacing to form a ULA. Consequently, any observed performance gap between the two configurations is attributable exclusively to spatial distribution gains, namely favorable propagation and macro-diversity, rather than any difference in transmit power. In particular, the centroid of the active antenna array remains at ≈ 3.5 m from the UE plane in both configurations and the per-UE distance to the array centroid varies by less than ± 0.7 dB of equivalent free-space path loss, so the proximity component of the comparison is held essentially constant. The 16 dB SINDR gap reported in Section VI therefore quantifies the architectural (favorable-propagation/macro-diversity) advantage of dMIMO net of any proximity effect, in line with the explicit decoupling described in Section VII-B.

Robustness of the bidirectional reciprocity-based calibration to rapid wireless-channel variation is built into the frame design: each frame is opened by a ZC preamble plus $KN_{r,k}$ orthogonal pilots and a dedicated equalization symbol (Section III), so the wireless channel and the effective UL response are re-estimated every frame and the calibration mapping the UL combiner to the DL precoder is applied to a channel at most one frame old. With a 110- μ s frame and an indoor 3.5 GHz coherence time of ~ 24 ms (Section VII-B, item (v)), per-frame ageing is two orders of magnitude inside the coherence time and the pilot-to-data ageing contribution is $\Delta \text{SINDR} \leq 0.4$ dB up to $v = 10$ m/s. The hardware mismatch matrices $\mathbf{D}_{TX}$, $\mathbf{D}_{RX,k}$ of (44) are channel-independent and drift only on thermal time scales, so the calibration retains its accuracy from indoor scenarios up to vehicular mobility.

VI. EXPERIMENTAL ANALYSIS

A. TIMING SYNCHRONIZATION ANALYSIS

For a configuration with two APs and four single-antenna UEs, Fig. 17 illustrates the received signal at each antenna of



FIGURE 16. Side view of the experimental setup. During measurements, all UEs are positioned to be equally spaced.




FIGURE 17. UL timing synchronization at AP 1 for four single-antenna UEs. The top row illustrates the received time-domain signals at each antenna, while the bottom row shows the corresponding cross-correlation outputs used for peak detection.

AP 1 in the UL: the top row shows the received time-domain signals and the bottom row the cross-correlation output, with sharp peaks at the frame-start samples and large peak-to-sidelobe ratios that confirm robust UL timing synchronization (similar observations apply to AP 2). The received signal appears circularly shifted, with portions split between the left/right boundaries of the observation window separated by a null interval (flat line) intentionally inserted at the transmitter as a visual sanity check; once the start sample is detected via the correlator of Section III, the samples are circularly shifted to restore the correct temporal order. Fig. 18 shows the corresponding DL received signal at each UE, again with distinct correlation peaks well above the sidelobes, confirming robust DL synchronization with the same circular-shift procedure.

B. CMIMO VERSUS DMIMO: FAVORABLE PROPAGATION ANALYSIS

For $N_{r,a} = 4$, $A = 2$, $K = 4$, $N_{r,k} = 1$, Figs. 19a and 19b report the worst-UE SINDR versus K for the four schemes (L4, L3,

L2 dMIMO using Algorithms 1–3, and a collocated $\lambda/2$ -ULA cMIMO baseline that itself uses L4 centralized processing as the upper bound of cMIMO performance), with all UEs collocated at $d_k \approx 15$ cm to enforce a worst-case inter-user correlation.

In the UL (Fig. 19a), the worst-UE SINDR decreases with K in every scheme but L4 dominates throughout: at $K = 2$ L4 reaches 28 dB versus 25 dB for cMIMO (3 dB gap), and at $K = 4$ cMIMO collapses to 8 dB while L4 holds 24 dB (16 dB gap). This is the favorable-propagation signature: geographically distributed APs see the collocated UEs from diverse angles and joint CPU-side processing orthogonalizes them, whereas a collocated ULA cannot resolve their spatial channels. L3 and L2 lie between L4 and cMIMO because their local-CSI precoders cannot fully suppress inter-UE interference under high correlation, but L2 still beats cMIMO by ≈ 5 dB at $K = 4$ thanks to AP geographic separation alone.

The DL (Fig. 19b) follows the same qualitative ordering, with two quantitative shifts driven by the absence of UL–DL




FIGURE 18. DL timing synchronization received at the UEs. The upper subfigures show the received time-domain waveforms containing null intervals, while the lower subfigures depict the cross-correlation peaks used to achieve DL synchronization.




FIGURE 19. Worst-UE SINDR for all processing levels (L4, L3, L2 dMIMO, and the cMIMO baseline). (a) UL SINDR vs. the number of UEs K : the centralized L4 architecture achieves a 16 dB gain over cMIMO at $K = 4$ due to spatial orthogonality. (b) DL SINDR vs. K : despite the 11 dB reciprocity-related drop relative to the UL, L4 dMIMO maintains a 15 dB advantage over cMIMO at $K = 3$. (c) UL under shadowing ($K = 2$): L4 demonstrates superior macro-diversity robustness, maintaining 30 dB versus 22 dB for cMIMO. (d) DL under shadowing ($K = 3$): all distributed architectures (L2, L3, L4) mitigate shadowing-induced losses more effectively than cMIMO despite the absence of hardware calibration.

hardware calibration in the duality-based Algorithms 1–3. First, every scheme loses 10–11 dB relative to its UL counterpart (e.g. L4: 35 dB UL vs. 24 dB DL at $K = 1$; cMIMO: 32 dB vs. 22 dB), exposing the reciprocity-mismatch term $\Delta \mathbf{H}$ in (44). Second, at $K = 3$ cMIMO collapses to 6 dB while L4 DL holds 22 dB, confirming that even without OTA calibration the favorable-propagation gain dwarfs the reciprocity loss in the dMIMO regime; L3 and L2 DL stay close to L4 because, as discussed in Section VI-C below, the reciprocity mismatch hits the globally aggregated channel of L4 harder than the locally aggregated channels of L3/L2. The take-away is that L4 centralized processing is the most effective approach to realize the dMIMO gains over cMIMO under the present hardware envelope.

C. CMIMO VERSUS DMIMO: MACRO-DIVERSITY ANALYSIS

In this subsection, we demonstrate the macro-diversity gains of dMIMO over cMIMO. To emulate a shadowing effect,

we place a metal sheet in front of the UEs and slide it horizontally from left to right. This creates a worst-case partial blockage scenario where the metal object progressively obstructs the line-of-sight path between the UEs and one of the APs. Figs. 19c and 19d illustrate the minimum per-user SINDR for both with-shadowing and without-shadowing conditions, evaluated across all considered processing schemes (L4, L3, L2, and cMIMO) in the UL (two users) and DL (three users), respectively.

In the UL (Fig. 19c), the no-shadowing trends mirror Fig. 19a at $K = 2$, but under shadowing L4 UL drops only from 32 dB to 30 dB while cMIMO drops from 29 dB to 22 dB, a 7 dB cMIMO loss versus a 2 dB L4 loss, leaving an 8 dB L4-over-cMIMO margin under shadowing that is the macro-diversity signature: when one AP is heavily shadowed, the remaining distributed APs sustain the link. L3 and L2 UL show the same qualitative robustness over cMIMO, confirming that geographic AP distribution mitigates localized shadowing through spatial diversity even without global combining.




FIGURE 20. Transmitted and received constellations in the high-SNR regime: (a) single-UE 4096-QAM (691.2 Mbps, 34.8 dB SINDR, BER = 0) and (b) two-UE 1024-QAM (1.152 Gbps aggregate, ≥ 31.5 dB per-UE SINDR).




FIGURE 21. Transmitted and received constellations for the four-user scenario using 256-QAM modulation (1.84 Gbps aggregate, ≥ 23.9 dB per-UE SINDR). All multi-UE results use the worst-case 15-cm UE spacing.

The DL (Fig. 19d) shows the same ordering, with cMIMO most affected by shadowing and L4/L3/L2 dMIMO maintaining substantially higher SINDR. Notably, the L4-versus-L3/L2 gap shrinks in the DL, especially at $K \in \{2, 3\}$: L4 DL precoding depends on the globally aggregated channel $\tilde{\mathbf{H}}_{k,c}$, which absorbs the full frequency-dependent TX–RX mismatch of (44), whereas L3/L2 local precoders use per-AP CSI and are inherently less sensitive to the global mismatch. The combined effect of shadowing plus uncalibrated reciprocity therefore favors L3/L2 in the DL more than in the UL, but in every case dMIMO outperforms cMIMO by a wide margin.

D. ROF PERFORMANCE IN THE HIGH SNR REGIME

Figs. 20a–21 demonstrate the SDRoF testbed’s peak-SNR capability through high-order QAM constellations, with the received signal coherently averaged over $N_{\text{avg}} = 200$ frames to lower the effective noise floor. The single-UE 4096-QAM case (Fig. 20a) achieves zero BER at 34.8 dB SINDR and 691.2 Mbps. The two-UE 1024-QAM case (Fig. 20b) achieves 31.5 dB minimum per-UE SINDR at 576 Mbps/UE (aggregate 1.152 Gbps). The four-UE 256-QAM case (Fig. 21) achieves 23.9 dB minimum per-UE SINDR at 460.8 Mbps/UE (aggregate 1.84 Gbps). These

are quasi-static peak-SNR figures rather than instantaneous real-time throughput, and all multi-UE results are taken with the worst-case ≈ 15 cm UE separation, so geographically distributed UE deployments would push the achievable peak even higher.

VII. DISCUSSION

A. PRACTICAL IMPLICATIONS AND THEORETICAL AGREEMENT

The experiments confirm that centralized L4 processing is required to fully exploit dMIMO’s spatial gains in dense multi-user scenarios: the 16 dB SINDR gain over cMIMO at $K = 4$ collocated UEs translates to a $\approx 2.7\times$ higher per-UE spectral efficiency under the Shannon bound, directly relevant to indoor hotspot densification. The 3 m AP separation used here is conservative relative to typical indoor offices (5–20 m), and because favorable-propagation gains grow with AP separation [3], [4], the 16 dB result is a lower bound on what is achievable in practice. Quantitatively, the experimentally observed gain matches the 10–20 dB SINDR gains predicted by theoretical CF studies for dense multi-user scenarios [2], [3], [4], providing the first direct wideband hardware validation of these predictions; the deviation from the upper bound reflects hardware impairments, the finite $N_t = 8$ active antennas, and indoor propagation effects

TABLE 3. Effective throughput derived from the measured SINDR of Figs. 19a–19d under 64-QAM rate-1/2 convolutional coding ($T_k \approx 122$ Mbps).

Scheme	Min. SINDR (dB)	K	Per-UE Tput (Mbps)	Aggregate Tput (Mbps)	Condition
L4 UL	24	4	122	488	no shadowing
L3 UL	16	4	122	488	no shadowing
L2 UL	13	4	122	488	no shadowing
cMIMO UL	8	4	—	122 [†]	no shadowing
L4 DL	22	3	122	366	no shadowing
cMIMO DL	6	3	—	0 [‡]	no shadowing
L4 UL	30	2	122	244	shadowing (one AP blocked)
cMIMO UL	22	2	122	244	shadowing, halved margin

[†] worst-UE-limited (cMIMO cannot serve four spatially-correlated UEs); [‡] decoding fails at 6 dB SINDR.

not captured by i.i.d. Rayleigh fading. The same theory [2], [3] predicts that distributed processing degrades more severely under highly correlated UE channels, which is precisely the collocated-UE worst case used in our setup, so under geographically distributed UEs the L4–L3/L2 gap is expected to narrow considerably (cf. Section VII-B). On the synchronization side, the 0.45 ps per-hop SDRoF jitter accumulation implies less than 3 dB SNR degradation across 20 cascaded hops at 3.5 GHz, confirming that building-scale SDRoF deployment is feasible without a dedicated synchronization infrastructure such as White Rabbit [39].

Beyond the indoor-hotspot use case, the wideband, low-jitter, bidirectional dMIMO/SDRoF testbed demonstrated here can also act as a physical-layer substrate for two emerging service classes. The first is multi-access edge computing for delay-sensitive workloads [40], [41], [42], [43]: the wireless link is usually the slowest and least reliable step in the device-edge-cloud offloading pipeline, and the 16 dB L4 SINDR gain and the calibration-free bidirectional operation reported in Section VI directly improve the per-link error and outage rate that the offloading and scheduling algorithms in those works rely on. The second is the digital-twin and human-digital-twin family of services [44], [45], [46], which run a continuous wireless link between a physical entity and its virtual model: the wideband (≈ 51 MHz) bidirectional channel of our testbed sustains the data rate a continuously updated twin needs, and the per-AP soft estimates produced by the L2 and L3 architectures expose enough statistical information for the federated-learning, semantic, and resource-allocation layers above to act on. A full algorithmic integration is outside the scope of this hardware-focused study, but the throughput, latency, and synchronization numbers reported here set the wireless-side limits within which the algorithms of the cited works must operate.

B. THROUGHPUT, AP SCALING, FAIR-COMPARISON METHODOLOGY, AND HARDWARE-IMPAIRMENT BOUNDS

1) THROUGHPUT

For convolutional coding ($r_c = 1/2$) with Viterbi decoding, the per-user net throughput at quasi-error-free SINDR is

$$T_k = \log_2(M) r_c \frac{C_a}{C + C_p} f_s, \quad (48)$$

evaluating to $T_k \approx 122$ Mbps for 64-QAM. At $K = 4$ collocated UEs, L4 UL therefore reaches a 488 Mbps aggregate while cMIMO is worst-UE limited to ≈ 122 Mbps; Table 3 summarizes the implied throughput for every scheme under both static and shadowing conditions. The high-SNR constellations of Section VI-D probe the maximum modulation order that the SDRoF chain can support without errors and yield uncoded raw line rates of 691.2 Mbps (4096-QAM, $K = 1$), 1.152 Gbps (1024-QAM, $K = 2$), and 1.84 Gbps (256-QAM, $K = 4$); the corresponding rate-1/2 net throughputs from (48) are 244, 407, and 651 Mbps, respectively.

2) SCALING BEYOND $A = 2$ APS AND MATCHED-DISTANCE METHODOLOGY

We chose the 2-AP setup so that the cMIMO/dMIMO comparison runs on identical hardware (same RRUs, same RF chains, same per-antenna power). Activating additional APs follows the standard scaling $\text{SINDR}_{L4} \propto N_t = A \cdot N_{t,a}$ [3], [4]; for $A = 4$ uncorrelated sites this yields an extra 3 dB, projecting UL/DL SINDR floors of $\approx 27/25$ dB at $K = 4$ collocated UEs, with cumulative SDRoF jitter still inside the 7 ps testbed bound (full derivation in Sec. III-B of the supplementary material). Translated to EVM through $\text{EVM} = 10^{-\text{SINDR}/20}$, the predicted worst-UE uplink EVM at $K = 4$ drops from $\approx 6.3\%$ at $A = 2$ (the present testbed) to $\approx 4.5\%$ at $A = 4$ and $\approx 3.2\%$ at $A = 8$. The 3GPP 256-QAM EVM requirement is 3.5%, so the system is expected to support 256-QAM once eight APs are active, even under the worst-case clustered-UE setup of Section VI. When the UEs are placed further apart in real deployments (Section VII-B, item (vi)), the EVM at the same A is expected to be lower still, because the L4 combining gain is no longer limited by the inter-user channel correlation. The dMIMO-vs-cMIMO comparison is constructed so that any proximity advantage is removed by design: (i) per-antenna transmit power is held at $P_{\max} = -3$ dBm in both configurations (total radiated power $P_{\max} + 10 \log_{10}(A N_{t,a}) = -3 + 9 = 6$ dBm across the eight active antennas); (ii) the same two RRUs (AP₁, AP₄) are used, with only the AP–AP spacing (3 m vs. $\lambda/2$) altered, while the array centroid stays at ≈ 3.5 m from the UE plane; and (iii) ZC pilot energy, frame structure, channel estimator and equalizer are identical. Under this matched-distance protocol, the per-UE path-loss difference

TABLE 4. Expected calibration-free UL–DL reciprocity SINDR gap as a function of carrier frequency on hardware similar to ours.

Band	f_c (GHz)	Gap (dB)	Source
n78 (FR1)	3.5	10–11	this work
n77 (FR1)	4.0	10–12	extrap. [24]
n79 (FR1)	5.5	11–12	extrap. [34]
n258 (FR2)	26.0	14–16	extrap. [24], [32]

between the two configurations is $\leq \pm 0.7$ dB, so the 16 dB gap is attributable to favorable propagation and macro-diversity rather than proximity. The shadowing experiment of Fig. 19c is an independent control because the metal sheet obstructs only one AP, leaving proximity unchanged but exposing the macro-diversity advantage in isolation.

3) HARDWARE-IMPAIRMENT BOUNDS

Where the current platform, fixed-tuned at 3.5 GHz, wall-mounted, 100 m of OM4 multi-mode fiber per AP, cannot be swept directly, we provide closed-form bounds derived from first-principles models in [11], [19], [20], [32], [34], and [47]. The bounds are summarized below; the full step-by-step derivations and supporting References appear in Sec. III-C through Sec. III-C through Sec. III-I of the supplementary material.

(i) *Power vs. OSR.* For the second-order ($L = 2$) SDM, doubling OSR yields $3(2L + 1) = 15$ dB of in-band SQNR for $2\times$ digital power, and the SDRoF advantage grows with bandwidth because DRoF linearity scales super-linearly [34], [36]. The RF band-pass filter at the RRU [37] is included in this budget; the SDM yields a substantial net hardware-cost reduction per RRU because it eliminates the multi-bit GSps DAC and the matched RF reconstruction chain that a comparable DRoF link would require.

(ii) *Carrier-frequency reciprocity gap.* Using the calibration-free RoF characterization in [24] and [34], the 10–11 dB UL–DL gap stays nearly constant inside FR1 and grows only at mmWave; the four-band prediction is summarized in Table 4. Specifically, the ± 1 dB FR1 spread (n77/n78/n79) means the present 3.5 GHz measurement is directly applicable to the dominant indoor-deployment band, while at $f_c > 24$ GHz the gap grows by ~ 4 –5 dB owing to stronger PA non-linearity and larger phase-noise variance, motivating OTA reciprocity calibration for any mmWave port.

(iii) *Long-term jitter.* The SDRoF synchronization architecture (PLL on the recovered QSFP clock [37], [38]) is phase-coherent by construction over an experimental session, so the 7 ps RMS sampling-clock jitter reported in Section VI-A is representative of operation over extended runs and does not impose a noticeable additional precoding penalty beyond what is already reflected in the measurement.

(iv) *Fronthaul length.* The end-to-end timing budget is dominated by the per-hop CDR jitter accumulating in quadrature, $\sigma_t \leq 0.45 \sqrt{N_h}$ ps (per-hop value measured on the present testbed), evaluating to ≈ 2.0 ps for $N_h = 20$

(multi-floor building) and ≈ 3.2 ps for $N_h = 50$ (campus-scale), both inside the 7 ps testbed bound. The TDD propagation budget at $\ell = 500$ m ($\approx 5\mu$ s round-trip) still fits within the 110- μ s OFDM frame and the 556-ns CP.

(v) *Mobility.* For UE motion at $v \leq 10$ m/s, the Doppler-induced ICI floor is bounded by ≤ -65 dB and the pilot-aging SINDR loss by ≤ 0.4 dB [11], [19], [47], so the $L4 > L3 > L2 > \text{cMIMO BER}$ hierarchy is preserved up to vehicular speeds.

(vi) *L4–L3 gap under inter-user correlation.* The CF analysis in [2] establishes that the L4–L3 SINDR gap shrinks monotonically as the average inter-user channel correlation $\bar{\rho}$ decreases. Our measurement under the clustered-UE configuration of Section VI (i.e., the high- $\bar{\rho}$ regime) shows an ~ 8 dB UL gap at $K = 4$ (Fig. 19a); under geographically separated UEs (low $\bar{\rho}$), the gap is expected to close substantially, in line with the theory. The shadowing experiment of Fig. 19c provides an independent control: it leaves UE separation unchanged but exposes the macro-diversity advantage, and L3/L2 retain robustness within a few dB of L4 despite using local-only precoders. The current testbed therefore gives a lower bound on the L2/L3 competitiveness in real deployments.

(vii) *Spatial-processing complexity.* The spatial processing in this work runs as a semi-real-time hardware-in-the-loop pipeline (a fully real-time setup would require porting the DSP onto the FPGA, beyond the scope of this experimental study): the central FPGA deinterleaves the SDM data and streams the per-subcarrier samples over PCIe to MATLAB on the host PC, where the spatial filtering and the per-frame combiner update for L2, L3, and L4 are executed. The main algorithmic cost is the matrix product $\mathbf{W}^H \mathbf{Y}$ plus the per-frame combiner update, of order $KL N_t$ complex multiply-accumulates per active subcarrier per OFDM symbol and an $N_t \times N_t$ generalized-eigenvalue solve. For the testbed parameters ($K = 4$, $L = 1$, $N_t = 8$, $C_a = 113$, OFDM symbol $\approx 110\mu$ s), this works out to about 33 MCMACs/s, well within the throughput of a single modern DSP/FPGA tile.

C. LIMITATIONS AND IMPLEMENTATION CAVEATS

For transparency, we summarize the limitations of the present testbed that bound the scope of the conclusions drawn in Sections VI-B.

Static channel. All measurements are performed under quasi-static indoor channels with stationary UEs, an artefact of the wall-mounted laboratory geometry. The mobility-induced loss is bounded analytically in item (v) above and remains below 0.4 dB up to $v = 10$ m/s, but the dynamic-channel BER hierarchy has not been measured.

Hardware reciprocity uncalibrated. The 10–11 dB UL–DL gap reported in Section VI reflects the calibration-free operating point: $\mathbf{D}_{\text{TX}}$ and $\mathbf{D}_{\text{RX},k}$ in (44) are not estimated and therefore enter as a residual term $\Delta \mathbf{H}_{k,c}$. OTA reciprocity calibration [24] is expected to recover most of this gap;

experimental confirmation requires dedicated calibration pilots that are not part of the present frame structure.

Single carrier frequency, single fronthaul length. The AAU mezzanine is fixed-tuned at 3.5 GHz with 100 m of OM4 multi-mode fiber per AP. The bounds in items (ii) and (iv) above predict that the present operating point is representative of the full FR1 band and of building-scale deployments up to $\ell \approx 500$ m, respectively, but the predictions have not been verified by frequency-swept or length-swept measurements on the present hardware.

Worst-case spatial correlation. The 15-cm UE clustering used throughout Section VI is well below the channel-decorrelation distance at 3.5 GHz; the SINDR gaps reported are lower bounds rather than typical values, and multi-room measurements with naturally separated UEs are expected to widen the cMIMO-versus-dMIMO margin and to narrow the L4-versus-L3/L2 gap (item (vi)).

Fixed coding rate and modulation. Convolutional rate-1/2 with Viterbi decoding is used uniformly; adaptive modulation and coding per UE would further increase aggregate throughput beyond the 1.84 Gbps reported, but is outside the scope of the present hardware-level study.

These caveats define the operating envelope of the reported results; the upcoming testbed iteration targets each of them, with the analytical bounds above and in Sec. III of the supplementary material serving as first-order predictions for the future measurements.

VIII. CONCLUSION

We presented a comprehensive wideband end-to-end experimental validation of a multi-user, multi-stream dMIMO testbed, comprising a CPU linked via SDRoF fronthaul to four APs and benchmarking the L4 (centralized), L3 (hybrid), and L2 (fully distributed) CF processing levels against a controlled cMIMO reference within a complete physical-layer framework (packet structure, pilot signaling, wideband OFDM, and UL–DL duality-based algorithms). The testbed reaches 4096-QAM, zero-BER, 1.86% EVM single-UE transmission and orthogonalizes four collocated UEs at 256-QAM with a 1.84 Gbps aggregate throughput, under daisy-chained SDRoF synchronization with only 0.45 ps jitter accumulation per hop. The headline result is that L4 processing yields a 16 dB UL SINDR gain over cMIMO at $K = 4$ collocated UEs and a 15 dB DL gain at $K = 3$ (cMIMO DL fails to decode beyond $K = 3$ at the calibration-free operating point), together with an 8 dB UL margin under shadowing (30 dB for L4 versus 22 dB for cMIMO), directly validating CF theoretical predictions while quantifying, for the first time, the 10–11 dB calibration-free UL–DL reciprocity penalty. The reported gains are a conservative lower bound: the experiments use a worst-case 3 m AP separation and worst-case 15 cm UE clustering, both of which understate the achievable performance in real deployments. Future work targets scaling to all four APs with geographically distributed mobile UEs, building-scale fronthaul lengths, carrier-frequency-swept reciprocity across

FR1/FR2, and OTA reciprocity calibration to compensate $\mathbf{D}_{\text{TX}}$ and $\mathbf{D}_{\text{RX},k}$ and recover the 10–11 dB DL gap, with the analytical bounds in Section VII-B (and the full derivations in Section III of the supplementary material) serving as first-order predictions for those measurements; integration of White Rabbit [39] and precision time protocol will be explored for fully real-time TDD at larger scale.

DATA AVAILABILITY

The experimental datasets and processing scripts used in this study are available from the corresponding author upon reasonable request.

ACKNOWLEDGMENT

A partial version of this work, specifically some part of Section IV and Figs. 20–21 from Section VI, has been accepted as a conference paper to the 2026 Joint European Conference on Networks and Communications & 6G Summit (EuCNC & 6G Summit).

REFERENCES

- [1] S. Elhoushy, M. Ibrahim, and W. Hamouda, “Cell-free massive MIMO: A survey,” *IEEE Commun. Surveys Tuts.*, vol. 24, no. 1, pp. 492–523, 1st Quart., 2022.
- [2] U. T. Demir, E. Björnson, and L. Sanguinetti, “Foundations of user-centric cell-free massive MIMO,” *Found. Trends Signal Process.*, vol. 14, nos. 3–4, pp. 162–472, Jan. 2021.
- [3] 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.
- [4] H. Q. Ngo, A. Ashikhmin, H. Yang, E. G. Larsson, and T. L. Marzetta, “Cell-free massive MIMO versus small cells,” *IEEE Trans. Wireless Commun.*, vol. 16, no. 3, pp. 1834–1850, Mar. 2017.
- [5] M. Munawar, M. Guenach, and I. Moerman, “Performance and architectural tradeoffs in scalable cell-free massive MIMO,” *IEEE Access*, vol. 12, pp. 150189–150203, 2024.
- [6] M. Munawar, M. Guenach, and I. Moerman, “On fairness degradation in cell-free multi-stream massive MIMO with mobile-grade antennas,” *IEEE Access*, vol. 14, pp. 21167–21180, 2026.
- [7] M. Farooq, H. Q. Ngo, E.-K. Hong, and L.-N. Tran, “Utility maximization for large-scale cell-free massive MIMO downlink,” *IEEE Trans. Commun.*, vol. 69, no. 10, pp. 7050–7062, Oct. 2021.
- [8] M. Bashar et al., “Exploiting deep learning in limited-fronthaul cell-free massive MIMO uplink,” *IEEE J. Sel. Areas Commun.*, vol. 38, no. 8, pp. 1678–1697, Aug. 2020.
- [9] M. Bashar, K. Cumanan, A. G. Burr, M. Debbah, and H. Q. Ngo, “On the uplink max–min SINR of cell-free massive MIMO systems,” *IEEE Trans. Wireless Commun.*, vol. 18, no. 4, pp. 2021–2036, Apr. 2019.
- [10] M. Chowdhury, J. Nam, and A. J. Goldsmith, “Coherence time of wireless channels with large antenna arrays,” in *Proc. IEEE Wireless Commun. Netw. Conf. (WCNC)*, San Francisco, CA, USA, Mar. 2017, pp. 1–6.
- [11] Y. Fang, L. Qiu, X. Liang, and C. Ren, “Cell-free massive MIMO systems with oscillator phase noise: Performance analysis and power control,” *IEEE Trans. Veh. Technol.*, vol. 70, no. 10, pp. 10048–10064, Oct. 2021.
- [12] Y. Zhang, Q. Zhang, H. Hu, L. Yang, and H. Zhu, “Cell-free massive MIMO systems with non-ideal hardware: Phase drifts and distortion noise,” *IEEE Trans. Veh. Technol.*, vol. 70, no. 11, pp. 11604–11618, Nov. 2021.
- [13] J. Zheng, J. Zhang, J. Cheng, V. C. M. Leung, D. W. K. Ng, and B. Ai, “Asynchronous cell-free massive MIMO with rate-splitting,” *IEEE J. Sel. Areas Commun.*, vol. 41, no. 5, pp. 1366–1382, May 2023.
- [14] E. Björnson and L. Sanguinetti, “Scalable cell-free massive MIMO systems,” *IEEE Trans. Commun.*, vol. 68, no. 7, pp. 4247–4261, Jul. 2020.
- [15] F. Fredj, Y. Al-Eryani, S. Maghsudi, M. Akrouit, and E. Hossain, “Distributed beamforming techniques for cell-free wireless networks using deep reinforcement learning,” *IEEE Trans. Cognit. Commun. Netw.*, vol. 8, no. 2, pp. 1186–1201, Jun. 2022.

- [16] M. Guenach, A. A. Gorji, and A. Bourdoux, "Joint power control and access point scheduling in fronthaul-constrained uplink cell-free massive MIMO systems," *IEEE Trans. Commun.*, vol. 69, no. 4, pp. 2709–2722, Apr. 2021.
- [17] M. Guenach, A. A. Gorji, and A. Bourdoux, "A deep neural architecture for real-time access point scheduling in uplink cell-free massive MIMO," *IEEE Trans. Wireless Commun.*, vol. 21, no. 3, pp. 1529–1541, Mar. 2022.
- [18] G. Femenias and F. Riera-Palou, "Fronthaul-constrained cell-free massive MIMO with low resolution ADCs," *IEEE Access*, vol. 8, pp. 116195–116215, 2020.
- [19] S. Kurma, K. Singh, P. K. Sharma, C.-P. Li, and T. A. Tsiftsis, "On the performance analysis of full-duplex cell-free massive MIMO with user mobility and imperfect CSI," *IEEE Trans. Commun.*, vol. 73, no. 5, pp. 3683–3701, May 2025.
- [20] B. Li, H. Zeng, X. Zhu, Y. Jiang, and Y. Wang, "Cooperative time synchronization and robust clock parameters estimation for time-sensitive cell-free massive MIMO systems," *IEEE Trans. Wireless Commun.*, vol. 23, no. 9, pp. 11552–11566, Sep. 2024.
- [21] E. Shi et al., "Joint AP-UE association and precoding for SIM-aided cell-free massive MIMO systems," *IEEE Trans. Wireless Commun.*, vol. 24, no. 6, pp. 5352–5367, Jun. 2025.
- [22] W. Xu, J. An, H. Li, L. Gan, and C. Yuen, "Algorithm-unrolling-based distributed optimization for RIS-assisted cell-free networks," *IEEE Internet Things J.*, vol. 11, no. 1, pp. 944–957, Jan. 2024.
- [23] D. Löschenbrand, M. Hofer, L. Bernadó, S. Zelenbaba, and T. Zemen, "Towards cell-free massive MIMO: A measurement-based analysis," *IEEE Access*, vol. 10, pp. 89232–89247, 2022.
- [24] Y. Cao et al., "Experimental performance evaluation of cell-free massive MIMO systems using COTS RRU with OTA reciprocity calibration and phase synchronization," *IEEE J. Sel. Areas Commun.*, vol. 41, no. 6, pp. 1620–1634, Jun. 2023.
- [25] J. R. Pérez, O. Fernández, L. Valle, A. Bedoui, M. Et-tolba, and R. P. Torres, "Experimental analysis of concentrated versus distributed massive MIMO in an indoor cell at 3.5 GHz," *Electronics*, vol. 10, no. 14, p. 1646, 2021.
- [26] P. Muralimohan, S. Bisoyi, P. R. Manne, H. K. Dureppagari, and K. Kuchi, "Experimental validation of large-scale distributed MIMO with real-time adaptive modulation," *IEEE Open J. Commun. Soc.*, vol. 7, pp. 1153–1162, 2026.
- [27] Y. Chu et al., "Testbed development: An intelligent O-RAN-based cell-free MIMO network," *IEEE Commun. Mag.*, vol. 63, no. 6, pp. 74–81, Jun. 2025.
- [28] M. Özden, C. Ülkü, and A. Özen, "An investigation of the effects of different UW sequences on the PAPR performance of UW-OFDM systems," in *Proc. 26th Signal Process. Commun. Appl. Conf. (SIU)*, May 2018, pp. 1–4.
- [29] S. Vishwanath, N. Jindal, and A. Goldsmith, "Duality, achievable rates, and sum-rate capacity of Gaussian MIMO broadcast channels," *IEEE Trans. Inf. Theory*, vol. 49, no. 10, pp. 2658–2668, Oct. 2003.
- [30] P. Viswanath and D. N. C. Tse, "Sum capacity of the vector Gaussian broadcast channel and uplink-downlink duality," *IEEE Trans. Inf. Theory*, vol. 49, no. 8, pp. 1912–1921, Aug. 2003.
- [31] M. Schubert and H. Boche, "Solution of the multiuser downlink beamforming problem with individual SINR constraints," *IEEE Trans. Veh. Technol.*, vol. 53, no. 1, pp. 18–28, Jan. 2004.
- [32] R. Schreier and G. C. Temes, *Understanding Delta-Sigma Data Converters*. Hoboken, NJ, USA: Wiley, 2004.
- [33] H. Li et al., "Real-time 100-GS/s sigma-delta modulator for all-digital radio-over-fiber transmission," *J. Lightw. Technol.*, vol. 38, no. 2, pp. 386–393, Jul. 29, 2019.
- [34] L. Breyne, G. Torfs, X. Yin, P. Demeester, and J. Bauwelinck, "Comparison between analog radio-over-fiber and sigma delta modulated radio-over-fiber," *IEEE Photon. Technol. Lett.*, vol. 29, no. 21, pp. 1808–1811, Nov. 1, 2017.
- [35] A. Checko et al., "Cloud RAN for mobile networks—A technology overview," *IEEE Commun. Surveys Tuts.*, vol. 17, no. 1, pp. 405–426, 1st Quart., 2015.
- [36] C. Ranaweera, E. Wong, A. Nirmalathas, C. Jayasundara, and C. Lim, "5G C-RAN with optical fronthaul: An analysis from a deployment perspective," *J. Lightw. Technol.*, vol. 36, no. 11, pp. 2059–2068, Jun. 1, 2018.
- [37] A. Vandierendonck et al., "A bi-directional distributed multi-user MIMO testbed using digital sigma-delta-over-fiber," *J. Lightw. Technol.*, vol. 43, no. 4, pp. 1595–1603, Feb. 1, 2025.
- [38] C.-Y. Wu, H. Li, J. Van Kerrebrouck, C. Meysmans, P. Demeester, and G. Torfs, "A bit-interleaved sigma-delta-over-fiber fronthaul network for frequency-synchronous distributed antenna systems," *Appl. Sci.*, vol. 11, no. 23, p. 11471, Dec. 2021.
- [39] P. Moreira, J. Serrano, T. Wlostowski, P. Loschmidt, and G. Gaderer, "White rabbit: Sub-nanosecond timing distribution over Ethernet," in *Proc. Int. Symp. Precis. Clock Synchronization Meas., Control Commun.*, Oct. 2009, pp. 1–5.
- [40] C. Yi, J. Cai, and Z. Su, "A multi-user mobile computation offloading and transmission scheduling mechanism for delay-sensitive applications," *IEEE Trans. Mobile Comput.*, vol. 19, no. 1, pp. 29–43, Jan. 2020.
- [41] K. Peng, B. Zhao, M. Bilal, and X. Xu, "Reliability-aware computation offloading for delay-sensitive applications in MEC-enabled aerial computing," *IEEE Trans. Green Commun. Netw.*, vol. 6, no. 3, pp. 1511–1519, Sep. 2022.
- [42] X. Ai, W. Liang, and C. Liu, "Joint optimization of model retraining and inference services in DT-assisted edge computing," *IEEE Trans. Netw.*, vol. 34, pp. 1804–1819, 2026.
- [43] B. Zhao, K. Peng, K. Zhang, H. Sun, Z. Tu, and D. Chu, "SerFlow: A multistage service-enhanced mechanism for workflow applications in CPSs with end-edge-cloud collaboration," *IEEE Internet Things J.*, vol. 12, no. 16, pp. 33792–33813, Aug. 2025.
- [44] S. D. Okegbile, J. Cai, H. Zheng, J. Chen, and C. Yi, "Differentially private federated multi-task learning framework for enhancing human-to-virtual connectivity in human digital twin," *IEEE J. Sel. Areas Commun.*, vol. 41, no. 11, pp. 3533–3547, Nov. 2023.
- [45] F. Zhu et al., "From data mirror to smart copilot: A survey on NextG semantic communication for propelling digital twin world into cognitive stage," *IEEE Commun. Surveys Tuts.*, vol. 28, pp. 4915–4947, 2026.
- [46] J. Chen et al., "Generative AI-aided QoE-aware resource allocations for RIS-assisted digital twin interaction with uncertain evolution," *IEEE Trans. Mobile Comput.*, vol. 25, no. 6, pp. 7888–7905, Jun. 2026.
- [47] J. Zheng, J. Zhang, E. Björnson, and B. Ai, "Impact of channel aging on cell-free massive MIMO over spatially correlated channels," *IEEE Trans. Wireless Commun.*, vol. 20, no. 10, pp. 6451–6466, Oct. 2021.



MUTEEN MUNAWAR (Student Member, IEEE) received the B.Sc. and M.Sc. degrees in electrical engineering in 2019 and 2021, respectively, and the second master's degree in electrical and information engineering from Seoul National University of Science and Technology, South Korea. He is currently pursuing the Ph.D. degree with UGent, Belgium. He was with the Communication and Signal Processing Laboratory, South Korea, from February 2023 to September 2023. Currently, he is working at IMEC, Leuven, as a Researcher. His research interests include wireless communications, applied machine learning, and digital multipliers.



ACHIM VANDIERENDONCK (Student Member, IEEE) received the M.S. degree in electrical engineering from Ghent University, Ghent, Belgium, in 2019. He is currently pursuing the Ph.D. degree with the Department of Information Technology, UGent. As part of the IDLab design group, he has designed and helped designing numerous chips related to high-speed analog and digital communication. His research interests include high-speed mixed-signal designs, digital circuit design, and clock and data recovery systems and radio-over-fiber wireless systems.



MAMOUN GUENACH (Senior Member, IEEE) received the Engineering degree from the École Mohammadia d'Ingénieurs, Morocco, in 1995, and the Ph.D. degree from UCLouvain, Belgium, in 2002. He was a Post-Doctoral Researcher with Ghent University from 2002 to 2006, where he has been a Visiting Research Professor since 2015. From 2006 to 2019, he was with Nokia Bell Laboratories. Since 2019, he has been a Research Scientist with imec, Leuven, Belgium. His research interests include coding, modulation, synchronization, and equalization for high-speed wired and wireless communications, including distributed massive MIMO.



CARO MEYSMANS (Student Member, IEEE) received the M.S. degree in electrical engineering from Ghent University, Ghent, Belgium, in 2020, where he is currently pursuing the Ph.D. degree with the Department of Information Technology. His research interests include analog circuit design, with a focus on time-domain processing and analog-to-digital converters and digital circuit design.



PARDISE ZARDOSHT received the M.S. degree in electrical engineering from Shiraz University, Shiraz, Iran, in 2017. She is currently pursuing the Ph.D. degree with the Department of Information Technology, UGent. Her research interests include digital and FPGA design for radio-over-fiber wireless systems and signal processing.



GUY TORFS (Senior Member, IEEE) received the Engineering degree in applied electronics and the Ph.D. degree in applied sciences and electronics from Ghent University, Ghent, Belgium, in 2007 and 2012, respectively.

Since 2011, he has been with IMEC, associated with Ghent University, where he first became a part-time Assistant Professor in 2015. In 2018, he became a full-time Associate Professor with Ghent University. His research interests include high-speed mixed-signal designs for wireless, fiber-optic and backplane communication systems, including digital signal processing and calibration, analog equalization circuits, radio-over-fiber transceivers, and clock and data recovery systems. In 2014, as part of the Bi-PON and Cascaded Bi-PON team, he was a recipient of the Greentouch 1000x Award. He was a co-recipient of a 2015 DesignCon Best Paper Award in the High-Speed Signal Design Category and the 2019 ECOC Best Demo Award. He has served as an Associate Editor for *IEICE Electronics Express* and *IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS—II: EXPRESS BRIEFS*, was a member of European Conference on Optical Communication (ECOC) TPC, and also a member of the IEEE International Solid-State Circuit Conference (ISSCC) TPC.



INGRID MOERMAN (Fellow, IEEE) received the degree in electrical engineering and the Ph.D. degree from Ghent University in 1987 and 1992, respectively. She became a part-time Professor with Ghent University in 2000. She is currently a Staff Member of IDLab, a core research group of IMEC, with research activities embedded at Ghent University and the University of Antwerp. She coordinates the research activities on intelligent wireless networking (iWiNe) at Ghent University, where she leads a team of more than 30 researchers. She is also the Program Manager of the Deterministic Networking Track, part of the Connectivity Program, IMEC. In this role, she coordinates research on end-to-end wired/wireless networking solutions for business and mission-critical applications that must meet strict quality of service requirements in terms of throughput, bounded latency, and reliability in private wireless environments. She has extensive experience in running and coordinating national and EU research-funded projects.