

Channel-Chart–Driven Antenna Activation for Energy-Efficient Cell-Free Massive MIMO

Yara Kafa, Luc Le Magoarou, Joumana Farah, Matthieu Crussière
Univ Rennes, INSA Rennes, CNRS, IETR-UMR 6164, Rennes, France

Abstract—Cell-free massive MIMO (CF-mMIMO) can provide uniformly high system performance across the coverage area through distributed cooperation. However, keeping all antennas active leads to high circuit and fronthaul power consumption. This paper studies energy-efficient uplink reception through *antenna activation*, where only a subset of distributed antennas is selected to serve the users. To enable network-wide activation with low signaling overhead, each access point (AP) maps its locally estimated channel state information (CSI) into a low-dimensional *channel-chart (CC) embedding* and forwards this representation to the central processing unit (CPU) instead of the full CSI. Based on the received embeddings, we formulate a binary antenna-activation problem for network EE maximization and propose three low-overhead strategies: (i) direct embedding-domain activation using a proxy EE metric without CSI reconstruction, (ii) nearest chart-point CSI approximation, which estimates CSI from a reference chart database for greedy EE optimization, and (iii) decoder-based CSI reconstruction using a pre-trained autoencoder. The proposed approaches further enable a systematic comparison of antenna activation techniques in terms of achieved performance, computational complexity, and signaling overhead. Numerical results show substantial EE gains over the all-active baseline, with only a minor degradation in sum spectral efficiency and a significant reduction in signaling overhead. Moreover, the proposed greedy framework achieves EE performance close to that of the simulated annealing baseline, confirming its effectiveness for low-complexity antenna activation.

Index Terms—Cell-free massive MIMO, Energy efficiency, Antenna activation, Channel charting, Low-dimensional CSI.

I. INTRODUCTION

CF-mMIMO (Cell-Free Massive MIMO) consists of a large set of geographically distributed antennas that serve users collaboratively, with their received uplink signals being coherently combined and processed at a central processing unit (CPU) [1], [2]. While this architecture provides strong macro-diversity and robustness, it requires operating a large number of radio-frequency (RF) chains and fronthaul links, leading to substantial circuit and fronthaul power consumption [3]. Consequently, keeping all distributed antennas active may yield only limited additional reception benefits while further increasing the network power cost. This motivates adaptive *network-wide antenna activation*, where only the most useful antennas are dynamically selected for uplink reception to improve energy efficiency while retaining the macro-diversity gains of CF-mMIMO. Energy-efficient antenna selection and activation have been widely studied in conventional MIMO, massive MIMO, and XL-MIMO systems [4]–[6], using channel-dependent selection rules based on instantaneous

CSI, imperfect CSI, long-term channel statistics, or learning-based optimization, and mainly focusing on centralized or co-located array deployments. In CF-mMIMO, related designs have mainly considered power control, AP selection, pilot assignment, and/or structured cooperation, generally relying on centralized optimization or coordination based on CSI-derived channel statistics [3], [7]. While effective, such approaches typically operate at the AP-UE association or co-operation level and do not address fine-grained, network-wide per-antenna activation from compact CSI representations, which is important for scalable dense deployments. Although joint antenna activation and power allocation under wireless fronthaul constraints has recently been investigated [8], that approach remains centralized and CSI/model-driven, requiring AP-UE access-channel statistics and wireless-fronthaul CSI to enforce SINR and fronthaul-capacity constraints. It therefore does not address antenna activation from compact learned CSI representations. This motivates *network-wide, explicit antenna activation* based on channel-charting (CC) embeddings, enabling centralized activation decisions with reduced signaling overhead. To address this gap, we exploit low-dimensional CC embeddings of locally estimated CSI [9]. Each AP, equipped with multiple antennas, compresses its instantaneous channel observations into a compact CC embedding, and forwards them to the CPU, thereby drastically reducing the fronthaul signaling overhead. The **contributions of this work** are as follows: we develop a greedy EE-maximization framework for scalable network-wide antenna activation under low signaling overhead, while closely approaching the performance of a full-CSI-based antenna-activation benchmark. Within this framework, we further propose three low-overhead activation strategies: a direct embedding-domain method that performs per-antenna activation across all APs directly in the compressed CC space by maximizing a proposed proxy EE metric, avoiding CSI reconstruction; a nearest chart-point CSI approximation strategy that estimates approximate CSI by matching the received CC embedding to a calibrated reference chart point and then performs greedy maximization of the exact EE objective; and a decoder-based strategy that instead reconstructs channel estimates from CC embeddings using a pre-trained autoencoder for greedy EE optimization. All methods are evaluated using realistic synthetic channels generated with the Sionna ray-tracing framework [10], where directive three-dimensional antenna radiation patterns with measured non-isotropic angular responses [11], [12] are integrated into the ray-tracing setup. The antennas are arranged in a cylindrical

array, producing directive beams that jointly cover the 360° azimuth range, thereby creating spatially selective channel conditions that support meaningful antenna-activation decisions across the distributed CF-mMIMO network. These methods enable a systematic comparison between direct compressed-domain activation and CSI-based activation with different reconstruction fidelities, thereby allowing us to evaluate the tradeoffs among activation performance, computational complexity, and required fronthaul signaling overhead. To the best of our knowledge, this is the first work to investigate network-wide, per-antenna activation in CF-mMIMO driven by CC representations.

II. SYSTEM MODEL

A. System Description

We consider the uplink of a user-centric CF-mMIMO system with L distributed APs serving K single-antenna users (UEs). Each AP $\ell \in \{1, \dots, L\}$ is equipped with N antennas, yielding a total of $M = LN$ antennas, and is connected to a CPU via fronthaul links. At each AP, the antennas are arranged in a uniform cylindrical array composed of N_{st} vertically stacked stages, each containing the same number of antennas uniformly distributed around the cylindrical aperture, so that the full array jointly covers the 360° azimuth range. The system operates in Time Division Duplexing (TDD) mode over a subband of bandwidth B_{sub} consisting of S subcarriers, where CSI is locally estimated at the APs using τ_p uplink pilot symbols. The effective uplink channel between user k and AP ℓ , incorporating the spatial selectivity induced by the directive three-dimensional antenna radiation patterns, is denoted by $\mathbf{h}_{k\ell} \in \mathbb{C}^N$, and the stacked global channel of user k is $\mathbf{h}_k \triangleq [\mathbf{h}_{k1}^\top \dots \mathbf{h}_{kL}^\top]^\top \in \mathbb{C}^M$. To limit signaling and processing complexity, user-centric dynamic cooperation clustering (DCC) at the AP level is adopted [7]. Let $\mathcal{M}_k$ and $\mathcal{A}_\ell$ denote the sets of APs serving user k and UEs served by AP ℓ , respectively, obtained using the AP-UE association algorithm from [7], with $|\mathcal{A}_\ell| \leq \tau_p$. The resulting association matrix $\mathbf{D} \in \{0, 1\}^{L \times K}$ has entries $D_{\ell k} = 1$ if AP ℓ serves user k , and 0 otherwise, and the corresponding DCC matrix for user k is $\mathbf{D}_k \triangleq \text{diag}(D_{1k}\mathbf{I}_N, \dots, D_{Lk}\mathbf{I}_N) \in \mathbb{C}^{M \times M}$.

B. Uplink Signal Model and Antenna Activation

The received uplink signal at AP ℓ is $\mathbf{r}_\ell = \sum_{k=1}^K \sqrt{p_k} \mathbf{h}_{k\ell} x_k + \mathbf{n}_\ell \in \mathbb{C}^N$, where x_k denotes the unit-power data signal of user k , p_k is the uplink transmit power, identical for all users, and $\mathbf{n}_\ell \sim \mathcal{CN}(\mathbf{0}, \sigma^2 \mathbf{I}_N)$ the additive noise. Antenna activation at AP ℓ is modeled by a binary vector $\boldsymbol{\alpha}_\ell \in \{0, 1\}^N$, whose n th entry $\alpha_{\ell,n} = 1$ if antenna n of AP ℓ is active and 0 otherwise. Collecting all per-AP activation vectors yields the global vector $\boldsymbol{\alpha} \triangleq [\boldsymbol{\alpha}_1^\top \dots \boldsymbol{\alpha}_L^\top]^\top \in \{0, 1\}^M$, and the corresponding global activation matrix $\boldsymbol{\Lambda} \triangleq \text{diag}(\boldsymbol{\alpha}) \in \{0, 1\}^{M \times M}$, with per-AP activation matrix $\boldsymbol{\Lambda}_\ell \triangleq \text{diag}(\boldsymbol{\alpha}_\ell)$.

C. Uplink Combining over Activated Antenna Ports

Once optimized at the CPU, the activation vector $\boldsymbol{\alpha}$ is fed back to the APs, where each AP applies its local matrix $\boldsymbol{\Lambda}_\ell$ so that uplink detection, baseband processing, and fronthaul forwarding involve only the selected antenna ports. MMSE-based combining is then adopted for uplink detection, considering both centralized and distributed implementations.

Centralized partial MMSE (C_MMSE): In the centralized scheme, the CPU performs uplink combining and thus requires the instantaneous CSI associated with the active antennas. To limit the combining complexity, partial MMSE (P-MMSE) reception is adopted by accounting only for the interference from users sharing at least one serving AP, as in [2]. Define $\mathcal{S}_k \triangleq \{i : \mathbf{D}_k \mathbf{D}_i \neq \mathbf{0}\}$. The combining vector for user k is

$$\mathbf{w}_k = \left(\sum_{i \in \mathcal{S}_k} p_i \mathbf{D}_k \boldsymbol{\Lambda}_i \mathbf{h}_i \mathbf{h}_i^H \boldsymbol{\Lambda}_k \mathbf{D}_k + \sigma^2 \mathbf{D}_k \boldsymbol{\Lambda}_k \mathbf{D}_k \right)^{-1} \mathbf{D}_k \boldsymbol{\Lambda}_k \mathbf{h}_k. \quad (1)$$

Distributed MMSE (D_MMSE): In the distributed scheme, each AP applies local MMSE detection, while the CPU performs equal-weight combining to the locally detected signals, as in [2]. The local combining vector for user k at AP ℓ is

$$\mathbf{w}_{k\ell} = \left(\sum_{i \in \mathcal{A}_\ell} p_i \boldsymbol{\Lambda}_\ell \mathbf{h}_{i\ell} \mathbf{h}_{i\ell}^H \boldsymbol{\Lambda}_\ell + \sigma^2 \boldsymbol{\Lambda}_\ell \right)^{-1} \boldsymbol{\Lambda}_\ell \mathbf{h}_{k\ell} \mathbf{D}_{\ell k}. \quad (2)$$

Unified Uplink SINR: Define the global combining vector $\mathbf{w}_k \in \mathbb{C}^M$ as the P-MMSE combiner in the centralized case, or as $\mathbf{w}_k \triangleq [\mathbf{w}_{k1}^\top \dots \mathbf{w}_{kL}^\top]^\top$ in the distributed case, with $\mathbf{w}_{k\ell} = \mathbf{0}$ for $\ell \notin \mathcal{M}_k$. The SINR of user k is

$$\gamma_k(\boldsymbol{\alpha}) = \frac{p_k |\mathbf{w}_k^H \mathbf{D}_k \boldsymbol{\Lambda}_k \mathbf{h}_k|^2}{\sum_{i \neq k} p_i |\mathbf{w}_k^H \mathbf{D}_k \boldsymbol{\Lambda}_i \mathbf{h}_i|^2 + \sigma^2 \mathbf{w}_k^H \mathbf{D}_k \boldsymbol{\Lambda}_k \mathbf{D}_k \mathbf{w}_k}. \quad (3)$$

III. ENERGY EFFICIENCY PROBLEM FORMULATION

We consider uplink EE maximization under antenna activation (AAc), formulated as

$$\underset{\boldsymbol{\alpha} \in \{0, 1\}^M}{\text{maximize}} \quad \text{EE}(\boldsymbol{\alpha}) \triangleq \frac{R(\boldsymbol{\alpha})}{P_{\text{tot}}(\boldsymbol{\alpha})}, \quad (4)$$

where the achievable uplink rate of user k is defined as $R_k(\boldsymbol{\alpha}) = B_{\text{sub}} \log_2(1 + \gamma_k(\boldsymbol{\alpha}))$. The corresponding uplink sum rate is $R(\boldsymbol{\alpha}) = \sum_{k=1}^K R_k(\boldsymbol{\alpha})$. AAc is constrained by the DCC association, such that antennas can be activated only at APs serving at least one user. Let $\delta_\ell \triangleq \mathbb{1}(\sum_{k=1}^K D_{\ell k} > 0)$ denote an indicator of whether AP ℓ serves at least one user. The total uplink power consumption is modeled following [13], adapted to the antenna activation context, as

$$P_{\text{tot}}(\boldsymbol{\alpha}) = \sum_{k=1}^K \frac{p_k}{\nu_k} + \sum_{\ell=1}^L \delta_\ell P_{\text{tc}} \sum_{n=1}^N \alpha_{\ell,n} + \sum_{\ell=1}^L P_{\text{fh},\ell}(\boldsymbol{\alpha}), \quad (5)$$

where $\nu_k \in (0, 1]$ is the power-amplifier efficiency at user k and P_{tc} is the internal circuit and processing power per active antenna. The fronthaul power consumption at AP ℓ

is $P_{\text{fh},\ell}(\alpha) = P_{\text{fix},\ell} + P_{\text{ft},\ell} \sum_{k=1}^K D_{\ell k} R_k(\alpha)$, where $P_{\text{fix},\ell}$ denotes the fixed fronthaul power and $P_{\text{ft},\ell}$ (W/(bit/s)) denotes the traffic-dependent power coefficient. Due to the binary activation variables and the coupling among antenna selection, multiuser interference, and MMSE combining, problem (4) is a nonconvex combinatorial fractional program. In the sequel, low-overhead CC-based strategies are developed to efficiently approximate solutions to (4).

IV. ANTENNA-ACTIVATION STRATEGIES UNDER LOW-OVERHEAD CSI EXCHANGE

This section presents three CC-based antenna-activation strategies for solving the EE-maximizing AAc problem in (4), whose objective is to determine the network-wide binary activation vector α . Instead of relying on full high-dimensional CSI at the CPU, these strategies use the channel-chart embeddings forwarded by the APs as input to the AAc procedure. This common CC-based formulation enables a systematic comparison among the three strategies in terms of the computational complexity required to compute α and the signaling overhead toward the CPU.

A. Channel-Chart Generation and Embedding Exchange

CC provides low-dimensional representations of locally estimated CSI that preserve neighborhood relationships (see [9], [14], [15] for the complete formulation). In this work, CC is used to enable AAc with reduced fronthaul signaling.

1) *Offline CC Construction:* During an offline calibration phase, each AP independently constructs a CC from locally estimated uplink CSI collected from a set of calibration users. Let $\mathbf{h}_\ell^{(i)} \in \mathbb{C}^N$ denote the effective channel of calibration user i at AP ℓ , obtained by averaging the per-subcarrier channel estimates over the S subcarriers in the considered subband. This averaging mitigates frequency-dependent variations and yields a geometrically robust representation $\mathbf{h}_\ell^{(i)}$, which is essential for learning a stable and meaningful CC. Following [9], the phase-insensitive similarity between calibration channels i and $j \neq i$ is defined as $s_\ell(i, j) \triangleq \frac{|\mathbf{h}_\ell^{(i)H} \mathbf{h}_\ell^{(j)}|}{\|\mathbf{h}_\ell^{(i)}\| \|\mathbf{h}_\ell^{(j)}\|} \in [0, 1]$, based on which the corresponding channel distance is given by $d_\ell(i, j) = \sqrt{2 - 2s_\ell(i, j)}$. Isomap, a nonlinear dimensionality reduction method [16], is then applied to the resulting pairwise distance matrix, whose entries contain the channel distances $d_\ell(i, j)$ between all calibration-user pairs $i \neq j$, to obtain the reference chart coordinates $\{\mathbf{z}_\ell^{(i)} \in \mathbb{R}^d\}$, where d denotes the CC dimension. This mapping places calibration users with similar channel characteristics at nearby chart points. The CC construction is performed offline and reused during online AAc operation.

2) *Forward Charting and Embedding Exchange:* During online operation, each AP maps newly observed CSI onto the precomputed CC through forward channel charting. Specifically, for user k served by AP ℓ , the instantaneous channel vector $\mathbf{h}_{k\ell}$ is compared with the stored calibration-user channels using the phase-insensitive similarity measure $s_\ell(k, i) \triangleq \frac{|\mathbf{h}_{k\ell}^H \mathbf{h}_\ell^{(i)}|}{\|\mathbf{h}_{k\ell}\| \|\mathbf{h}_\ell^{(i)}\|}$, for $i = 1, \dots, N_c$, where N_c denotes

the number of calibration users at AP ℓ . The K_{nn} most similar calibration-user samples are selected, and their corresponding chart coordinates are combined through weighted interpolation to obtain the low-dimensional embedding $\mathbf{z}_{k\ell} \in \mathbb{R}^d$. This feed-forward procedure [9], [15] has linear complexity in the number of calibration users and relies on only local CSI. The resulting embeddings $\mathbf{z}_{k\ell}$ are then sent to the CPU, enabling centralized AAc with substantially reduced fronthaul signaling compared with full CSI exchange.

B. Antenna-Activation Strategies Based on Channel Charts

Given the CC embeddings available at the CPU, we propose three low-overhead AAc strategies for solving problem (4), each exploiting the embeddings in a different way.

1) Direct Embedding-Domain Antenna Activation (DD):

This strategy performs AAc in the CC domain without reconstructing instantaneous CSI at the CPU. The embeddings are instead used to define a proxy EE metric that guides antenna selection with minimal computational and signaling overhead.

a) *Embedding-to-Channel-Power Prediction:* The approach relies on the observation that users with correlated channel vectors at an AP exhibit similar per-antenna channel power profiles. Since the CC is constructed to preserve this local similarity between neighboring users, proximity in the embedding space enables channel-power prediction without requiring instantaneous CSI at the CPU. During the offline phase, each AP ℓ constructs a CPU-stored calibration lookup table, $\{(\mathbf{z}_\ell^{(i)}, \mathbf{u}_\ell^{(i)})\}_{i=1}^{N_c}$, where i indexes calibration users at AP ℓ , $\mathbf{z}_\ell^{(i)} \in \mathbb{R}^d$ is a calibration-user embedding, and $\mathbf{u}_\ell^{(i)} = [|h_{\ell,1}^{(i)}|^2, \dots, |h_{\ell,N}^{(i)}|^2]^\top$ is the corresponding per-antenna channel power vector. During online operation, for a user k with embedding $\mathbf{z}_{k\ell}$, the CPU searches this lookup table to identify the set $\mathcal{N}_{k\ell}$ of nearest calibration users. The predicted per-antenna channel power vector is obtained as $\hat{\mathbf{u}}_{k\ell} = \sum_{i \in \mathcal{N}_{k\ell}} \omega_{k\ell}^{(i)} \mathbf{u}_\ell^{(i)}$, where the weights are given by

$$\omega_{k\ell}^{(i)} = \frac{\|\mathbf{z}_{k\ell} - \mathbf{z}_\ell^{(i)}\|_2^{-1}}{\sum_{j \in \mathcal{N}_{k\ell}} \|\mathbf{z}_{k\ell} - \mathbf{z}_\ell^{(j)}\|_2^{-1}}. \quad (6)$$

Stacking across APs yields the predicted channel-power matrix $\hat{\mathbf{U}} \in \mathbb{R}_+^{K \times M}$.

b) *Embedding-Based User-Similarity Metric:* During inference, the average similarity of user k at AP ℓ is computed as $\rho_{k\ell} = \frac{1}{K-1} \sum_{i \neq k} (1 - \frac{1}{2} \|\mathbf{z}_{k\ell} - \mathbf{z}_{i\ell}\|_2^2)$, where $\|\mathbf{z}_{k\ell} - \mathbf{z}_{i\ell}\|_2$ denotes the normalized Euclidean distance between the chart embeddings of users k and i . This metric reflects how closely user k is clustered with the other users in the chart space at AP ℓ and is later used through the factor $(1 - \rho_{k\ell})$ to penalize such contributions in the proxy utility.

c) *Embedding-Based Proxy SINR and EE:* Let $M_{\text{act}}(\alpha) = \sum_{\ell=1}^L \sum_{n=1}^N \alpha_{\ell,n}$ denote the total number of activated antennas. Using the predicted channel-power matrix $\hat{\mathbf{U}}$ and the embedding-based similarity metric $\rho_{k\ell}$, we define the proxy utility of user k as

$$\hat{\eta}_k(\alpha) = \sum_{\ell=1}^L \sum_{n=1}^N \alpha_{\ell,n} \hat{u}_{k\ell,n} \frac{1 - \rho_{k\ell}}{1 + \sum_{n'=1}^N \alpha_{\ell,n'}}, \quad (7)$$

where $\sum_{n'=1}^N \alpha_{\ell, n'}$ denotes the number of antennas activated at AP ℓ under the candidate activation vector α . Each term in (7) quantifies the desirability of activating antenna n at AP ℓ for user k . It rewards antennas with high predicted channel power, penalizes AP-user contributions associated with highly clustered users through the factor $(1 - \rho_{k\ell})$, and promotes sparse activation at each AP through the normalization by the number of active antennas at that AP. Based on $\hat{\eta}_k(\alpha)$, the proxy SINR is defined as

$$\hat{\gamma}_k(\alpha) = \frac{p_k}{\sigma^2} \hat{\eta}_k(\alpha) \frac{M_{\text{act}}(\alpha) - K}{M_{\text{act}}(\alpha)}, \quad M_{\text{act}}(\alpha) > K. \quad (8)$$

This proxy captures the effects of the transmit signal-to-noise ratio of user k , the defined proxy utility, and the activated-antenna factor. The term $\frac{M_{\text{act}}(\alpha) - K}{M_{\text{act}}(\alpha)}$ favors activation sets with more antennas than users, reflecting the additional spatial degrees of freedom available for multi-antenna reception.

The proxy $\hat{\text{EE}}(\alpha)$ is obtained by replacing the exact SINR $\gamma_k(\alpha)$ in (4) with the proxy SINR $\hat{\gamma}_k(\alpha)$.

d) *Greedy Antenna Activation Using Proxy EE*: AAc is performed using Algorithm 1. A low-overhead solution is obtained by replacing the true $\text{EE}(\alpha)$ with the proxy $\hat{\text{EE}}(\alpha)$, computed from CC-derived quantities via (7) and (8). The candidate and active antenna sets are denoted by $\mathcal{C}$ and $\mathcal{S}$, respectively. The algorithm uses a warm start by activating one antenna per user based on the predicted channel powers $\hat{u}_{k,m}$, and then incrementally expands $\mathcal{S}$ to maximize $\hat{\text{EE}}(\alpha)$. Leveraging the empirically observed concave dependence of EE on the number of active antennas, the procedure terminates when no further improvement is achieved.

2) *Nearest Chart-Point CSI Approximation (NN)*: This strategy uses CC embeddings to obtain approximate instantaneous CSI at the CPU and performs AAc by evaluating the EE objective in (4) using the resulting channel approximation. In an offline calibration phase, each AP ℓ collects local CSI realizations $\{\mathbf{h}_\ell^{(i)}\}$ from calibration users, computes the associated chart coordinates $\{\mathbf{z}_\ell^{(i)}\}$, and provides the resulting CSI-embedding pairs to the CPU as a lookup table. During online operation, given the embedding $\mathbf{z}_{k\ell}$ of user k , the CPU retrieves the nearest calibration-user embedding following a CC-based nearest-neighbor CSI retrieval method similar to [17], i.e., $i_{k\ell}^* = \arg \min_i \|\mathbf{z}_{k\ell} - \mathbf{z}_\ell^{(i)}\|_2$, and reuses the corresponding channel realization as $\hat{\mathbf{h}}_{k\ell} = \mathbf{h}_\ell^{(i_{k\ell}^*)}$. Stacking $\{\hat{\mathbf{h}}_{k\ell}\}$ over all users and APs yields the approximate global CSI matrix $\hat{\mathbf{H}}$, and the antenna-activation vector is obtained by running Algorithm 1. In this strategy, $\hat{\mathbf{H}}$ is used both to compute the warm-start channel powers $u_{k,m} = |\hat{h}_{k,m}|^2$ and to evaluate the objective $\text{EE}(\alpha)$ during the greedy search.

3) *Decoder-Based CSI Reconstruction (RC)*: Following the CC-based CSI compression and reconstruction framework in [18], a third strategy employs a neural encoder-decoder to reconstruct the instantaneous CSI at the CPU from low-dimensional CC embeddings, which is then used for AAc. For each served user k , AP ℓ compresses its locally estimated subband-averaged effective channel $\mathbf{h}_{k\ell} \in \mathbb{C}^N$ into an em-

Algorithm 1 Greedy antenna activation

Input: $\{\mathbf{z}_{k\ell}\}_{k,\ell}$, $\{p_k\}$, $\mathbf{D}$
Output: α^* , M_{act}^*
Init: $\mathcal{C} = \{1, \dots, LN\}$, $\mathcal{S} \leftarrow \emptyset$, $\alpha \leftarrow \mathbf{0}$, $M_{\text{act}} \leftarrow K + 1$, $\text{EE}_{\text{max}} \leftarrow -\infty$, $\alpha^* \leftarrow \mathbf{0}$
1: **for** $k = 1, \dots, K$ **do** ▷ Warm start
2: $m^* \leftarrow \arg \max_{m \in \mathcal{C}} u_{k,m} // u_{k,m} = |h_{k,m}|^2$
3: **if** $\alpha_{m^*} = 0$ **then**
4: $\alpha_{m^*} \leftarrow 1$, $\mathcal{S} \leftarrow \mathcal{S} \cup \{m^*\}$
5: **end if**
6: **end for**
7: **while** true **do**
8: **while** $|\mathcal{S}| < M_{\text{act}}$ **do**
9: $m^* \leftarrow \arg \max_{m \in \mathcal{C} \setminus \mathcal{S}} \text{EE}(\alpha | \alpha_m = 1)$
10: $\alpha_{m^*} \leftarrow 1$, $\mathcal{S} \leftarrow \mathcal{S} \cup \{m^*\}$
11: **end while**
12: $\text{EE}_{\text{cur}} \leftarrow \text{EE}(\alpha)$
13: **if** $\text{EE}_{\text{cur}} > \text{EE}_{\text{max}}$ **then**
14: $\text{EE}_{\text{max}} \leftarrow \text{EE}_{\text{cur}}$, $\alpha^* \leftarrow \alpha$, $M_{\text{act}}^* \leftarrow M_{\text{act}}$
15: **else**
16: **break**
17: **end if**
18: $M_{\text{act}} \leftarrow M_{\text{act}} + 1$
19: **end while**
20: **return** α^* , M_{act}^*

bedding $\mathbf{z}_{k\ell} \in \mathbb{R}^d$ and forwards it to the CPU. Using an AP-specific decoder $g_\ell(\cdot)$ trained offline as in [18], the CPU maps $\mathbf{z}_{k\ell}$ to a normalized channel direction $\tilde{\mathbf{h}}_{k\ell} = g_\ell(\mathbf{z}_{k\ell}) \in \mathbb{C}^N$ with $\|\tilde{\mathbf{h}}_{k\ell}\|_2 = 1$. Since this normalized output preserves only the channel direction and discards the magnitude information, AP ℓ additionally forwards the channel norm $\beta_{k\ell} = \|\mathbf{h}_{k\ell}\|_2$ as side information, allowing the CPU to restore the large-scale channel strength. In addition to this magnitude information, phase side information is required because CC encoding relies on phase-insensitive features, which may cause the reconstructed channels to exhibit AP-dependent phase ambiguities. Inspired by the phase-alignment mechanism in [18], we adopt a lightweight approximation tailored to the uniform cylindrical array with N_{st} vertical stages. Specifically, to provide a coarse phase reference for reconstruction, each AP computes a circular-mean phase $\phi_{k\ell,s}$ for each stage s from the locally estimated CSI and forwards these phase values as additional side information, i.e., $\{\phi_{k\ell,s}\}_{s=1}^{N_{\text{st}}}$ along with $(\mathbf{z}_{k\ell}, \beta_{k\ell})$. The CPU then reconstructs each antenna coefficient as $\hat{h}_{k\ell,n} = \beta_{k\ell} \tilde{h}_{k\ell,n} e^{j\phi_{k\ell,n}}$. Stacking the reconstructed vectors across all APs yields the global reconstructed CSI $\hat{\mathbf{h}}_k \in \mathbb{C}^M$ for user k . The corresponding user-wise vectors are then collected into $\hat{\mathbf{H}}_{\text{recon}}$, and the antenna-activation vector is obtained by running Algorithm 1. In this strategy, $\hat{\mathbf{H}}_{\text{recon}}$ is used both to compute the warm-start channel powers $u_{k,m} = |\hat{h}_{k,m}|^2$ and to evaluate the objective $\text{EE}(\alpha)$ during the greedy search.

TABLE I: Uplink fronthaul overhead per UE-AP link

Method	AAc Overhead	Total Overhead (AAc + Dist. Det.)	Total Overhead (AAc + Cent. Det.)
Full CSI	$2NS$	$2NS$	$2NS$
DD	d	d	$d + 2N_{\ell}^{\text{act}}S$
NN	d	d	$d + 2N_{\ell}^{\text{act}}S$
RC	$d + N_{\text{st}} + 1$	$d + N_{\text{st}} + 1$	$d + N_{\text{st}} + 1 + 2N_{\ell}^{\text{act}}S$

TABLE II: Complexity summary of the proposed AAc strategies

Method	CSI/Channel-Power Acquisition	Objective Evaluation
DD	$\mathcal{O}(LKN_c d)$	$\mathcal{O}(KM_{\text{act}})$
NN	$\mathcal{O}(LKN_c d)$	$\mathcal{O}(KM_{\text{act}}^3)$
RC	$\mathcal{O}(LK(FT + T^2 + TN))$	$\mathcal{O}(KM_{\text{act}}^3)$

C. Fronthaul Overhead and Computational Complexity

Before the numerical results, we compare the proposed CC-based AAc strategies in terms of fronthaul signaling toward the CPU and online CSI/channel-power acquisition and EE-evaluation cost.

1) *Fronthaul Overhead Comparison:* Table I reports the uplink fronthaul overhead per UE-AP link, measured as the number of transmitted real scalars, for AAc and uplink detection. Obtaining an AAc solution based on full instantaneous CSI at the CPU requires forwarding $2NS$ real scalars, whereas DD and NN require only the d -dimensional CC embedding, and RC requires $d + N_{\text{st}} + 1$ scalars due to the additional channel-norm and phase-alignment information, making DD and NN the most compact representations. For distributed detection, no additional CSI is exchanged beyond that required for AAc; hence, the total fronthaul coincides with the activation overhead, as reported in the “AAc + Dist. Det.” column. For centralized detection, AP ℓ additionally forwards the CSI of its active antennas, adding $2N_{\ell}^{\text{act}}S$ real scalars to the AAc signaling cost, where N_{ℓ}^{act} is the number of active antennas at AP ℓ , as reported in the “AAc + Cent. Det.” column.

2) *Computational Complexity Analysis:* We compare the AAc strategies in terms of CSI/channel-power acquisition and objective-evaluation cost, as summarized in Table II. DD and NN rely on a CPU-stored lookup table and perform nearest-neighbor search over N_c calibration embeddings per AP, yielding an online acquisition complexity of $\mathcal{O}(LKN_c d)$; this step predicts per-antenna channel-power vectors for DD and retrieves approximate CSI for NN. RC avoids lookup-table search by reconstructing CSI with a trained decoder, with online cost $\mathcal{O}(LK(FT + T^2 + TN))$, where F and T are decoder parameters denoting the number of random Fourier features and the hidden-layer width, respectively. The DD strategy evaluates a CC-based proxy EE with complexity $\mathcal{O}(KM_{\text{act}})$ per objective evaluation, while NN and RC rely on the exact MMSE-based EE, whose dominant cost is the MMSE covariance-matrix inversion, scaling as $\mathcal{O}(KM_{\text{act}}^3)$ per objective evaluation. Overall, DD has the lowest total online computational complexity among the three strategies, owing to its embedding-domain operation and low-complexity proxy-EE evaluation.

TABLE III: Simulation Parameters

Param.	Value	Param.	Value
(f_c, B_{sub}) [MHz]	(3800, 0.36)	(S, τ_p)	(12, 15)
(L, N, K, N_{st})	(10, 48, 30, 4)	d	4
(p_k, σ^2) [dBm]	(10, -93)	ν_k	0.4
$(P_{\text{fix}}, P_{\text{tc}})$ [W]	(0.825, 0.2)	P_{ft} [W/bps]	0.25×10^{-9}
F	200	T	128

V. PERFORMANCE EVALUATION

This section evaluates the DD, NN, and RC strategies. These methods are compared with ground-truth (GT) benchmarks, corresponding to solutions obtained under the availability of full instantaneous CSI at the CPU. Each curve is labeled following the convention strategy-combiner-algorithm, where the strategy is one of the proposed schemes (DD, NN, RC) or the GT benchmark, the combiner is centralized (C_MMSE) or distributed (D_MMSE) MMSE detection, and the algorithm indicates how the activation vector is obtained: greedy (G), simulated annealing (SA), or all-active (AA). The greedy algorithm (G) refers to Algorithm 1; when paired with the GT strategy, it operates on full CSI at the CPU. Simulated annealing (SA), implemented following [19], uses a search procedure different from the proposed greedy algorithm and directly evaluates the exact EE objective. The all-active case (AA) keeps all antennas active, applying no activation-search algorithm. The remainder of this section presents the simulation setup and evaluates the proposed strategies and benchmarks in terms of EE and sum spectral efficiency (SE), concluding with a summary of the resulting performance trends.

Simulation Setup: The proposed framework is evaluated using realistic datasets generated with *Sionna* [10] for the *Étoile* urban scenario in Paris, with the main parameters summarized in Table III. For DD and NN, $N_c = 8000$ calibration realizations are used for CC construction and 2000 realizations for testing, while RC uses $N_c = 3500, 4500$, and 2000 realizations for CC calibration, autoencoder training, and testing, respectively.

Performance comparisons: Fig. 1 shows the empirical CDFs of EE and sum SE over 1000 Monte-Carlo realizations. The all-active baseline GT-C_MMSE-AA achieves the highest SE but the lowest EE, due to excessive circuit and processing power consumption, resulting in a left-shifted EE curve with limited variability. In contrast, the evaluated AAc schemes substantially improve EE, shifting the EE CDFs to the right while incurring only moderate SE losses. The proposed CC-driven greedy methods, DD-C_MMSE-G, NN-C_MMSE-G, and RC-C_MMSE-G, achieve average EE gains of approximately 280-315% over GT-C_MMSE-AA while activating only 40-45 antennas out of $M = 480$. Among them, NN-C_MMSE-G is closest to the full-CSI-based GT benchmarks, whereas DD-C_MMSE-G and RC-C_MMSE-G are slightly degraded by proxy-based optimization and decoder reconstruction loss, respectively. GT-C_MMSE-SA achieves the highest observed EE, while GT-C_MMSE-G yields a nearly identical CDF at significantly lower complexity. In our implementation,




Fig. 1. CDFs of (top) EE [Mbit/Joule] and (bottom) Sum spectral efficiency [bit/s/Hz] for the considered AAC strategies under MMSE reception.

EE performance	DD	<	RC	<	NN
Signaling efficiency	RC	<	DD	=	NN
Computational efficiency	RC	<	NN	<	DD
	better $\rightarrow$				

Fig. 2. Qualitative EE-signaling-complexity trade-off among the proposed CC-driven AAC strategies.

SA performs $N_{SA} = 1000$ candidate-swap trials per tested antenna cardinality, each requiring one exact EE evaluation with complexity $\mathcal{O}(KM_{act}^3)$; further implementation details are omitted due to space limitations, and a general description of SA can be found in [19]. By contrast, GT-D_MMSE-SA shows a pronounced EE and SE degradation compared with centralized detection, confirming that centralized detection is important for high performance, while the proposed CC-driven strategies aim to make it scalable by reducing CSI signaling overhead.

Trade-off summary: Fig. 2 summarizes the performance-overhead-complexity trade-offs among the proposed AAC strategies. NN achieves the best EE performance and the strongest performance-overhead compromise, making it the preferred strategy under the considered setup. DD provides the lowest signaling and computational cost at the price of lower EE accuracy, whereas RC improves over DD in terms of EE but requires additional side information and higher decoder-based reconstruction complexity.

VI. CONCLUSION

This work demonstrated that channel charting enables scalable centralized antenna activation in CF-mMIMO systems by replacing full instantaneous CSI exchange with compact CC embeddings. Three CC-based strategies—DD, NN, and RC—were proposed to exploit these embeddings for network-wide EE optimization under different levels of CSI approximation, revealing clear trade-offs between performance, signaling, and

computational complexity. The results, obtained using realistic antenna radiation patterns in a ray-tracing environment, show that the proposed strategies achieve competitive EE performance with respect to the considered full-CSI benchmarks while requiring substantially lower signaling. These findings confirm that CC-based representations can successfully support centralized energy-efficient antenna activation, with the preferred strategy depending on whether the main priority is signaling overhead, online complexity, or EE performance.

Acknowledgment: This work was supported by the ANR under the France 2030 program, grant “NF-PERSEUS: ANR-22-PEFT-0004”.

REFERENCES

- [1] 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, 2017.
- [2] U. T. Demir, E. Björnson, and L. Sanguinetti, “Foundations of User-Centric Cell-Free Massive MIMO,” *Found. Trends Signal Process.*, vol. 14, no. 3–4, pp. 162–472, 2021.
- [3] H. Q. Ngo, L.-N. Tran, T. Q. Duong, M. Matthaiou, and E. G. Larsson, “On the Total Energy Efficiency of Cell-Free Massive MIMO,” *IEEE Trans. Green Commun. Netw.*, vol. 2, no. 1, pp. 25–39, 2018.
- [4] N. P. Le, F. Safaei, and L. C. Tran, “Antenna Selection Strategies for MIMO-OFDM Wireless Systems: An Energy Efficiency Perspective,” *IEEE Trans. Veh. Technol.*, vol. 65, no. 4, pp. 2048–2062, 2016.
- [5] M. Guo and M. C. Gursoy, “Statistical Learning Based Joint Antenna Selection and User Scheduling for Single-Cell Massive MIMO Systems,” *IEEE Trans. Green Commun. Netw.*, vol. 5, no. 1, pp. 471–483, 2021.
- [6] J. C. Marinello, T. Abrão, A. Amiri *et al.*, “Antenna Selection for Improving Energy Efficiency in XL-MIMO Systems,” *IEEE Trans. Veh. Technol.*, vol. 69, no. 11, pp. 13 305–13 318, 2020.
- [7] S. Chen, J. Zhang, E. Björnson, J. Zhang, and B. Ai, “Structured Massive Access for Scalable Cell-Free Massive MIMO Systems,” *IEEE J. Sel. Areas Commun.*, vol. 39, no. 4, pp. 1086–1100, 2021.
- [8] O. A. Topal, O. T. Demir, E. Björnson, and C. Cavdar, “Energy-Efficient Cell-Free Massive MIMO with Wireless Fronthaul,” in *Proc. Asilomar Conf. Signals, Syst., Comput.*, 2024, pp. 1591–1596.
- [9] L. L. Magoarou, “Efficient Channel Charting via Phase-Insensitive Distance Computation,” *IEEE Wireless Commun. Lett.*, vol. 10, no. 12, pp. 2634–2638, 2021.
- [10] J. Hoydis, S. Cammerer, F. Ait Aoudia *et al.*, “Sionna,” 2022, <https://nvlabs.github.io/sionna/>.
- [11] L. Batel, A. Clemente, and C. Delaveaud, “Superdirective and Compact Electronically-Beam-Switchable Antenna for Smart Communication Objects,” in *Proc. Eur. Conf. Antennas Propag. (EuCAP)*, 2019, pp. 1–4.
- [12] A. Tornese, A. Clemente, and C. Delaveaud, “Maximizing Gain of Endfire Arrays Based on Electrical Dipoles Using Spherical Harmonics,” *IEEE Trans. Antennas Propag.*, vol. 72, no. 10, pp. 7525–7535, 2024.
- [13] N. Ghiasi *et al.*, “Energy Efficient AP Selection for Cell-Free Massive MIMO Systems: Deep Reinforcement Learning Approach,” *IEEE Trans. Green Commun. Netw.*, vol. 7, no. 1, pp. 29–41, 2023.
- [14] C. Studer, S. Medjkouh, E. Gonultas *et al.*, “Channel Charting: Locating Users Within the Radio Environment Using Channel State Information,” *IEEE Access*, vol. 6, pp. 47 682–47 698, 2018.
- [15] L. L. Magoarou, T. Yassine, S. Paquelet, and M. Crussière, “Channel Charting Based Beamforming,” in *Proc. Asilomar Conf. Signals, Syst., Comput.*, 2022, pp. 1185–1189.
- [16] J. B. Tenenbaum, V. de Silva, and J. C. Langford, “A Global Geometric Framework for Nonlinear Dimensionality Reduction,” *Science*, vol. 290, no. 5500, pp. 2319–2323, 2000.
- [17] P. Stephan, F. Euchner, and S. T. Brink, “Channel Charting-Based Channel Prediction on Real-World Distributed Massive MIMO CSI,” in *Proc. IEEE WCNC*, 2025, pp. 1–6.
- [18] H. Abbass, L. L. Magoarou, J. Farah, and M. Crussière, “Channel Charting for Efficient Precoding in Cell-Free Massive MIMO,” in *Proc. Asilomar Conf. Signals, Syst., Comput.*, 2025, pp. 868–874.
- [19] D. Delahaye, S. Chaimatanan, and M. Mongeau, *Simulated Annealing: From Basics to Applications*. Cham: Springer International Publishing, 2019, pp. 1–35.