

Dynamic antenna selection for hybrid beamforming in ISAC systems: a joint transceiver optimization framework

Chenghao Zhu, Ziyi Gong, Liang Wu^(✉), Zaichen Zhang, and Jiangzhou Wang

Abstract Integrated sensing and communication (ISAC) is a key enabler for 6G networks, yet its performance is critically constrained by the trade-off between hardware complexity and system flexibility. Conventional hybrid beamforming (HBF) architectures suffer from rigid antenna connections, limiting their adaptability in dynamic environments such as unmanned aerial vehicle (UAV) networks. To overcome this, we propose a dynamic antenna selection (AS) framework for HBF-based ISAC systems with joint transceiver optimization. The proposed architecture adaptively reconfigures both transmit and receive antenna subsets, forming a closed-loop system where high-precision sensing continuously guides robust communication. A weighted objective function is formulated to maximize both spectral efficiency and sensing accuracy (characterized by the Cramér–Rao bound). To solve this mixed-integer non-convex problem, we introduce a Gumbel-Softmax relaxation that transforms the combinatorial AS task into a differentiable form, enabling gradient-based end-to-end optimization within an alternating optimization framework. Furthermore, an iterative adaptive approach (IAA) is employed for robust single-snapshot angle estimation under the dynamically formed non-uniform receive array. Simulation results demonstrate that the proposed joint transceiver optimization significantly outperforms both transmitter-only and full-array baselines, achieving a superior communication-sensing trade-off and establishing a new Pareto frontier for AS-aided ISAC systems.

Keywords Integrated sensing and communication, hybrid beamforming, dynamic antenna selection, joint transceiver optimization, Gumbel-Softmax.

1 Introduction

Integrated sensing and communication (ISAC) is a cornerstone of 6G networks, enabling the co-design of sensing and communication on shared hardware and spectrum to improve spectral and energy efficiency [1–5]. Emerging applications such as intelligent transportation and unmanned aerial vehicle (UAV) networks impose stringent demands on both high-resolution sensing and reliable multi-user communication in highly dynamic environments [6, 7]. Beamforming plays a central role in this integration, allowing simultaneous optimization of multi-user links and high-resolution sensing [8–10]. However, the performance of ISAC in massive multiple-input-multiple-output (MIMO) systems is fundamentally constrained by a trade-off among performance, hardware cost, and power consumption [11], making the choice of beamforming architecture particularly critical.

This trade-off has led to several beamforming architectures with different balances of performance and hardware complexity. Fully digital beamforming (DBF), while delivering optimal performance, requires one radio frequency (RF) chain per antenna, making its hardware cost prohibitive for large-scale arrays [12]. At the other extreme, fully analog beamforming (ABF) uses only one RF chain but supports only a single data stream, severely limiting its multi-user and multi-target capabilities [13]. Hybrid beamforming (HBF) strikes a pragmatic balance by combining a low-dimensional digital precoder with an analog phase-shifter network, achieving near-digital performance with significantly reduced hardware complexity [14]. Recent studies have demonstrated HBF's effectiveness in ISAC systems, including alternating minimization for joint beamforming [15] and sum-rate maximization in full-duplex scenarios [16]. However, a critical limitation persists in conventional HBF architectures: each RF chain is permanently connected to a fixed subset of antennas through a static subarray configuration. This fixed connection pattern locks the available spatial degrees of freedom (DoFs) into a predetermined structure, preventing the system from adapting to time-varying channel conditions and target dynamics [17].

This structural rigidity manifests itself in three fundamental limitations that critically constrain the performance of HBF-

• Chenghao Zhu is with the National Mobile Communications Research Laboratory, Southeast University, Nanjing 210096, China. E-mail: chenghaozhu@seu.edu.cn.

• Ziyi Gong is with Purple Mountain Laboratories, Nanjing 211111, China. E-mail: gongziyi@pmlabs.com.cn.

• Liang Wu, Zaichen Zhang, and Jiangzhou Wang are with the National Mobile Communications Research Laboratory, Southeast University, Nanjing 210096, China, and also with Purple Mountain Laboratories, Nanjing 211111, China. E-mail: wuliang@seu.edu.cn; zczhang@seu.edu.cn; j.z.wang@seu.edu.cn.

Manuscript received: ; revised: ; accepted:

based ISAC systems in dynamic UAV scenarios.

First, the inability to adapt to task dynamics [18]. In highly mobile environments, the instantaneous channel conditions dictate a time-varying optimal subset of antennas for transmission and reception. A static subarray configuration, however, permanently binds each RF chain to a fixed group of antennas, preventing the system from selecting the globally optimal antenna set that aligns with the transient channel state. This fundamental mismatch results in diminished beamforming gain and less effective interference suppression.

Second, the lack of hardware elasticity [19]. Practical systems must contend with component failures, energy-saving requirements, and varying traffic loads. A fixed subarray architecture cannot gracefully degrade or reconfigure without manual intervention. The ability to dynamically activate or deactivate individual antennas would enable fault-tolerant operation and adaptive power management.

Third, suboptimal beam quality [17]. The effective array aperture and beam pattern are locked by the fixed subarray geometry. Dynamic antenna geometry can reshape the virtual array by choosing the most favorable antenna positions, thereby improving both communication directivity and sensing angular resolution.

To overcome these limitations, dynamic antenna selection (AS) has emerged as a promising technique that enables the activation of the most effective antenna subsets [20]. By allowing each RF chain to dynamically connect to the best antennas for the current channel state, AS effectively restores the spatial degrees of freedom that are otherwise locked by fixed subarrays. However, a critical review of existing AS literature reveals notable deficiencies that hinder its direct application to ISAC. The majority of research [21] is confined to communication-only contexts, employing methods that cannot be readily extended to the dual-objective optimization of communication and sensing performance required by ISAC. Furthermore, many schemes [22–24] restrict AS to the transmitter only, neglecting the receiver side and thus failing to unlock full end-to-end performance. Other approaches either suffer from the prohibitive computational complexity of combinatorial searches or, as seen in deep learning solutions [25, 26], rely on idealized feedback and idealized objective functions lacking ISAC-specific considerations.

To address these challenges, this paper proposes a novel framework for joint transceiver antenna selection and hybrid beamforming in ISAC systems, specifically designed for low-altitude UAV scenarios. The proposed architecture features a dynamic, AS-aided HBF structure that adaptively reconfigures both the transmit and receive antenna subsets. Unlike existing

works that optimize the transmitter and receiver separately, our framework establishes a closed-loop system where high-precision sensing information is continuously fed back to guide the transmitter's antenna selection and beamforming strategy. This sensing-aided communication loop enables the system to align its spatial resources with the instantaneous environment, directly tackling the rigidity of fixed subarrays.

It is worth noting that although low-altitude UAV channels may experience blockage, multipath propagation, and Doppler-induced fast time-varying fading in practice, this work focuses on the dominant line-of-sight (LoS) propagation regime that commonly exists between a UAV and an elevated base station operating at millimeter-wave frequencies [27, 28]. This modeling choice is intentionally adopted to isolate the impact of the proposed dynamic antenna selection and hybrid beamforming framework on the communication–sensing trade-off from channel estimation errors. The proposed closed-loop architecture is not restricted to the LoS model: it only requires periodically updated channel and angle estimates, and therefore can be naturally extended to more practical channel models through the feedback mechanism discussed in Section 2.

The core technical challenge lies in the joint optimization of the binary antenna selection matrices and the continuous beamforming matrices, which leads to a highly complex mixed-integer non-convex problem. To overcome this, we develop a differentiable optimization framework based on the Gumbel-Softmax relaxation. This technique transforms the intractable combinatorial selection task into a continuous form, enabling gradient-based end-to-end optimization of the selection schemes alongside the analog and digital beamformers. The resulting optimization is efficiently solved within an alternating optimization (AO) framework that iteratively updates each variable block. Furthermore, to handle the non-uniform receive array geometry induced by receiver-side antenna selection, we employ the iterative adaptive approach (IAA) for robust single-snapshot angle estimation, which closes the sensing loop by providing accurate angular information for the next transmission cycle.

The main contributions of this paper are summarized as follows, each directly addressing one of the limitations discussed above.

- To break the rigidity of fixed subarrays in dynamic environments, we propose a dynamic AS-aided HBF architecture that adaptively reconfigures both the transmit and receive antenna subsets. By establishing a closed-loop system where high-precision sensing feedback continuously guides the transmitter's configuration, the pro-

posed framework enables the system to align its spatial resources with instantaneous channel conditions, directly overcoming the inability of fixed subarrays to adapt to task dynamics.

- To overcome the combinatorial complexity of joint AS and beamforming optimization, we introduce a Gumbel-Softmax relaxation that transforms the intractable binary selection problem into a continuous, differentiable form. This enables gradient-based end-to-end optimization of the selection matrices alongside the analog and digital beamformers, avoiding exhaustive search or heuristic methods. The resulting alternating optimization framework is provably convergent and achieves polynomial complexity, addressing the scalability concerns of existing approaches.
- To explicitly manage the trade-off between communication and sensing, we formulate a weighted objective function of spectral efficiency (SE) and inverse Cramér-Rao bound (CRB). The weighting factor provides a tunable knob for prioritizing either function according to application requirements, enabling a graceful and controllable balance that is essential for practical ISAC deployment.
- To enable robust angle estimation under the non-uniform receive array induced by receiver-side AS, we employ the IAA for single-snapshot angle estimation. The IAA algorithm achieves near-CRB performance without requiring uniform array assumptions or multiple snapshots, thereby closing the sensing loop with high precision and supporting real-time tracking in dynamic UAV environments.

The remainder of this paper is organized as follows. In Section 2, we introduce the system model for the AS-aided HBF ISAC system and formulate the joint transceiver optimization problem. Section 3 details the proposed differentiable relaxation technique, which leverages the Gumbel-Softmax method to handle the binary nature of the AS. In Section 4, we develop an alternating optimization framework to solve the formulated problem by iteratively updating the beamformers and selection matrices. Section 5 presents the IAA algorithm for robust angle estimation with the non-uniform array. Simulation results are provided in Section 6 to validate the performance of our proposed framework, and finally, the paper is concluded in Section 7.

Notations: The bold lower-case and upper-case letter denote a vector and a matrix, respectively. $(\cdot)^T$ and $(\cdot)^H$ denote the transpose and conjugate transpose operations, respectively. $\odot$ denotes the Hadamard product. $\mathbf{I}_N$ denotes the




Fig. 1 The ISAC scenario for low-altitude UAVs.

N -dimensional identity matrix. $\mathbb{C}^{M \times N}$ denotes the set of $M \times N$ dimensional complex matrices. $\text{diag}(a_1, \dots, a_N)$ denotes a diagonal matrix composed of $a_1, \dots, a_N$. $\text{Re}\{\cdot\}$ denotes the real part of a value. $\mathbb{E}[\cdot]$ denotes the expectation. $\text{Tr}(\cdot)$ denotes the trace of a matrix. $|\cdot|$ denotes the modulus of a value. $\|\cdot\|_F$ denotes the Frobenius norm of a matrix. a_i denotes the i -th element of the vector $\mathbf{a}$. $[\mathbf{A}]_{i,j}$ and $[\mathbf{A}]_{:,j}$ denotes the (i, j) -th element and the j -th column of the matrix $\mathbf{A}$, respectively. $\nabla_x y$ denotes the gradient of y with respect to x . $\frac{\partial y}{\partial x}$ denotes the partial derivative of y with respect to x .

2 System Model and Problem Formulation

Consider a low-altitude UAV scenario, where K UAVs with single antenna are served by a dual-functional radar-communication (DFRC) BS [29], shown in Fig. 1. The BS is equipped with N_{TX} transmit antennas and N_{RF} RF chains, and each RF chain is connected to every antenna through a phase shifter (PS) and a switch. The dimension of the transmitted signal is equal to the number of UAVs. Based on the above structure, the analog beamformer implemented by the PS network is expressed as $\mathbf{F}_{\text{RF}} \in \mathbb{C}^{N_{\text{TX}} \times N_{\text{RF}}}$, where $[\mathbf{F}_{\text{RF}}]_{m,n} = e^{j\phi_{m,n}}$ with $\phi_{m,n}$ denoting the shifted phase of the m -th antenna connected to the n -th RF chain. The digital baseband beamformer is expressed as $\mathbf{F}_{\text{BB}} \in \mathbb{C}^{N_{\text{RF}} \times K}$ and the transmitted signal is expressed as $\mathbf{s} \in \mathbb{C}^{K \times 1}$, satisfying $\mathbb{E}[\mathbf{s}\mathbf{s}^H] = \mathbf{I}_K$.

The proposed framework establishes a closed-loop operation for ISAC, as illustrated in Fig. 2. The process is initialized with a coarse angle estimate (e.g., from a wider discovery beam or prior tracking data [30]) and iteratively refines both




Fig. 2 The AS-aided transceiver architecture and the closed-loop ISAC framework. The operation consists of three steps: (1) Adaptive transmission based on prior knowledge; (2) Echo reception and angle estimation; and (3) Feedback of updated angles to refine the next transmission.

the angle estimates and the transceiver configuration through the following three steps:

- **Step 1: Adaptive Transmission.** At the beginning of each new closed-loop cycle, the BS uses the updated angle estimates obtained from Step 3 of the previous cycle as prior spatial information. Based on these refreshed estimates, the transmit antenna selection matrix and hybrid beamformers are jointly optimized before the next ISAC transmission. The ISAC signals are then transmitted via the AS-aided HBF architecture, where each RF chain activates a specific subset of antennas to form desired beams towards the designated UAVs.
- **Step 2: Echo Reception & Sensing.** In the sensing phase, the switch network dynamically reconfigures the receive array. By connecting each RF chain to a unique antenna, the system effectively forms a sparse, fully-digital array. The BS receives the reflected echoes to update the angular information.
- **Step 3: Feedback & Loop Closure.** Finally, the newly acquired high-precision angular estimates are fed back to the transmitter side. This information dictates the transceiver optimization for the next transmission, thereby realizing a robust sensing-aided communication loop.

Remark 1: The RF switches used for antenna selection have typical switching times on the order of nanoseconds to microseconds [31], while the channel coherence time in UAV scenarios at 30 GHz is on the order of 0.1–several

milliseconds [32]. Thus the switching overhead is 2–5 orders of magnitude smaller than the coherence time and can be safely neglected.

Remark 2: The proposed closed-loop framework does not rely on perfectly known angle information. As described above, the process is initialized with only a coarse angle estimate, (e.g., obtained from a wider discovery beam or prior tracking data), which may contain estimation errors. The target angles are then refined by the IAA-based sensing module from the received echoes and are periodically updated within each channel coherence interval. These refreshed angle estimates are fed back to the transmitter to reconstruct the antenna selection and the beamformers for the next transmission cycle. Therefore, neither the initialization nor the subsequent optimization requires perfectly known angle information. The complex channel gain is assumed known in the simulation to isolate the algorithm design from channel estimation errors. In practice, it can be estimated from pilot signals due to its slow variation.

To examine whether the initial angle estimation error could accumulate, we consider the closed-loop tracking simulation in Section 6.5 (Fig. 11). The initial angle estimates are assumed to contain a moderate error, yet the closed-loop configuration still maintains high sensing accuracy and spectral efficiency before the first sudden movement. Moreover, even after each sudden target movement, which causes a much larger angle mismatch, the system re-converges rapidly. These results confirm that angle estimation errors do not accumulate and that the initial estimation error does not pose a fundamental limitation.

The proposed framework operates in a periodic tracking-and-transmission manner. During each channel coherence interval, the IAA-based sensing module first estimates the target angles from the received echoes. These estimates are then fed back to the transmitter, and the alternating optimization algorithm described in Section 4 is executed to update both the antenna selection matrices and the hybrid beamformers before the next transmission interval. Since RF switch reconfiguration is significantly faster than the channel coherence time (Remark 1), the entire update can be completed within one coherence block. Consequently, the proposed closed-loop architecture can naturally track slowly varying UAV trajectories through repeated optimization cycles while maintaining computational feasibility in dynamic scenarios.

The HBF architecture incorporates a dynamic AS mechanism. We define the binary selection matrix for the transmit array as $\mathbf{\Omega} \in \{0, 1\}^{N_{\text{TX}} \times N_{\text{RF}}}$, where $[\mathbf{\Omega}]_{i,j} = 1$ indicates that the i -th antenna is connected to the j -th RF chain. Con-

sequently, the effective analog beamformer is given by the element-wise product $\tilde{\mathbf{F}}_{\text{RF}} = \mathbf{\Omega} \odot \mathbf{F}_{\text{RF}} \in \mathbb{C}^{N_{\text{TX}} \times N_{\text{RF}}}$. Assuming a LoS channel model, the channel vector between the BS and the k -th UAV is modeled as $\mathbf{h}_k = \gamma_k \mathbf{a}(\theta_k) \in \mathbb{C}^{N_{\text{TX}} \times 1}$, where $\mathbf{a}(\theta_k) \in \mathbb{C}^{N_{\text{TX}} \times 1}$ is the array steering vector corresponding to the k -th UAV. γ_k is the complex channel gain accounting for the large-scale fading and is given by $\gamma_k = \frac{\lambda}{4\pi d_k} e^{-j\frac{2\pi d_k}{\lambda}}$, where λ is the carrier wavelength and d_k is the distance between the BS and the k -th UAV. The achievable SE for the k -th UAV is thus expressed as

$$R_k = \log_2 \left(1 + \frac{|\mathbf{h}_k^H \tilde{\mathbf{F}}_{\text{RF}} \mathbf{v}_k|^2}{\sum_{l \neq k} |\mathbf{h}_k^H \tilde{\mathbf{F}}_{\text{RF}} \mathbf{v}_l|^2 + \sigma_k^2} \right), \quad (1)$$

where $\mathbf{v}_k = [\mathbf{F}_{\text{BB}}]_{:,k}$ is the digital beamforming vector for the k -th UAV's data stream, and σ_k^2 is the variance of the additive white Gaussian noise (AWGN) at the k -th UAV receiver.

For sensing, the BS receives the echoes reflected from the UAVs. Similar to the transmitter, the receive array employs an AS strategy, represented by the binary selection vector $\omega \in \{0, 1\}^{N_{\text{TX}} \times 1}$. The received echo signal is given by

$$\mathbf{r} = \sum_{k=1}^K \beta_k \mathbf{b}(\theta_k) \mathbf{a}^H(\theta_k) \tilde{\mathbf{F}}_{\text{RF}} \mathbf{F}_{\text{BBS}} + \tilde{\mathbf{n}}, \quad (2)$$

where $\mathbf{b}(\theta_k) = \omega \odot \mathbf{a}(\theta_k)$ is the effective steering vector. β_k represents the complex round-trip channel coefficient of the k -th UAV, which incorporates both the radar cross-section (RCS) and the round-trip path loss given by $\beta_k = \frac{\lambda}{(4\pi)^{1.5} d_k^2} \sqrt{\sigma_{\text{RCS},k}} e^{-j\frac{4\pi d_k}{\lambda}}$, where $\sigma_{\text{RCS},k}$ denotes the RCS of the k -th UAV. The noise term $\tilde{\mathbf{n}} = \omega \odot \mathbf{n}$ represents the sampled noise, with $\mathbf{n} \sim \mathcal{CN}(0, \sigma_s^2 \mathbf{I}_{N_{\text{TX}}})$ being the original complex Gaussian noise vector.

To quantify the sensing performance, we adopt the CRB for the angle estimation, which provides a lower bound on the variance of any unbiased estimator. As derived in Appendix A, the CRB for estimating θ_k is given by

$$\begin{aligned} \text{CRB}(\theta_k) = & \frac{\sigma_s^2}{2|\beta_k|^2} \left[\dot{\mathbf{b}}^H(\theta_k) \dot{\mathbf{b}}(\theta_k) \sum_{l=1}^K |\mathbf{a}^H(\theta_k) \tilde{\mathbf{F}}_{\text{RF}} \mathbf{v}_l|^2 \right. \\ & \left. + \mathbf{b}^H(\theta_k) \mathbf{b}(\theta_k) \sum_{l=1}^K |\dot{\mathbf{a}}^H(\theta_k) \tilde{\mathbf{F}}_{\text{RF}} \mathbf{v}_l|^2 \right]^{-1}, \end{aligned} \quad (3)$$

where $\dot{\mathbf{a}}(\theta_k) = \frac{\partial \mathbf{a}(\theta_k)}{\partial \theta_k}$ and $\dot{\mathbf{b}}(\theta_k) = \frac{\partial \mathbf{b}(\theta_k)}{\partial \theta_k}$ denote the first-order derivatives of the respective steering vectors with respect to θ_k .

Our objective is to jointly optimize the transceiver design to maximize both communication throughput and sensing accuracy. We formulate this multi-objective problem by max-

imizing a weighted sum of the total communication SE and the total inverse CRB, which characterizes the sensing precision. To balance the numerical scales of these two metrics, a weighting factor ρ is introduced. The resulting optimization problem is formulated as

$$\underset{\mathbf{F}_{\text{RF}}, \mathbf{F}_{\text{BB}}, \mathbf{\Omega}, \omega}{\text{maximize}} \quad \sum_{k=1}^K R_k + \rho \cdot \text{CRB}^{-1}(\theta_k) \quad (4a)$$

$$\text{subject to} \quad |\mathbf{F}_{\text{RF}}]_{i,j}| = 1, \forall i, j, \quad (4b)$$

$$\|\tilde{\mathbf{F}}_{\text{RF}} \mathbf{F}_{\text{BB}}\|_F^2 \leq P_t, \quad (4c)$$

$$\mathbf{\Omega} \in \{0, 1\}^{N_{\text{TX}} \times N_{\text{RF}}}, \quad (4d)$$

$$\omega \in \{0, 1\}^{N_{\text{TX}} \times 1}, \quad (4e)$$

$$\sum_{i=1}^{N_{\text{TX}}} \omega_i = N_{\text{RF}}. \quad (4f)$$

Here, (4b) is the unit-modulus constraint imposed by the phase shifters. Constraint (4c) limits the total transmit power to P_t . The constraints (4d) and (4e) enforce the binary nature of the antenna selection at both the transmitter and receiver. The core challenge of this problem lies in the non-convex objective function and the combinatorial constraints induced by the integer variables $\mathbf{\Omega}$ and ω .

3 Differentiable Relaxation of Antenna Selection

The primary bottleneck in solving the optimization problem (4) lies in the binary selection matrices, $\mathbf{\Omega}$ and ω . Their discrete nature creates a non-differentiable optimization landscape, which prevents the use of efficient gradient-based methods by impeding gradient flow from the objective function back to the selection parameters. To circumvent this, the proposed approach is to construct a differentiable proxy for these binary variables. This enables all components of the transceiver, including the digital and analog beamformers as well as the transmit and receive AS schemes, to be jointly optimized in an end-to-end manner using gradient-based methods.

The most straightforward implementation of such a proxy is to relax the binary variables into continuous parameters within the $[0, 1]$ interval. However, this naive relaxation is fundamentally flawed. During the optimization procedure, without a robust mechanism to enforce decisive, near-binary selections, these parameters often collapse towards ambiguous intermediate values around 0.5. This outcome represents a state of maximum uncertainty from the optimizer's perspective, and when these values are inevitably quantized back to 0 or 1 post-optimization, the resulting discrete selection leads to a significantly suboptimal performance. This

highlights the need for a more sophisticated relaxation technique that encourages near-binary outputs while maintaining differentiability.

To address the problem of ambiguous intermediate values, we adopt a reparameterization approach based on the Gumbel-Softmax trick [33, 34]. Instead of directly learning continuous parameters in $[0, 1]$, the network learns a set of unconstrained parameters known as logits [35], which represent the unnormalized log-probabilities for each categorical choice. For each antenna, we have a binary choice, that is, to select it or not. This corresponds to a categorical distribution with $N = 2$ categories.

Let the logits for these two choices be α_1 (select) and α_2 (not select). We add Gumbel noise to these logits. This step is crucial for two interconnected reasons. Firstly, it introduces controlled stochasticity into the selection process, encouraging the model to explore different antenna configurations rather than prematurely converging to a suboptimal local minimum. Without this noise, the model's output would be deterministic, hindering its ability to discover more effective selection policies. Secondly, the addition of Gumbel noise is the foundation of the reparameterization trick that allows us to approximate the non-differentiable sampling process with a continuous and differentiable function, enabling end-to-end training with gradient-based methods.

By applying the Softmax function, a differentiable, “soft” sample is obtained as

$$z_1 = \frac{\exp((\alpha_1 + g_1)/\tau)}{\sum_n \exp((\alpha_n + g_n)/\tau)}, \quad (5)$$

where z_1 represents the “soft” probability of selecting the antenna, $g_n \sim \text{Gumbel}(0, 1)$ is random noise, generated via inverse transform sampling: $g_n = -\log(-\log(u_n))$, $u_n \sim U(0, 1)$, and τ is the temperature. A smaller τ forces the output z_1 to be closer to 0 or 1, thus encouraging a definitive selection and preventing convergence to 0.5.

For a binary case ($N = 2$), this formulation can be simplified. Since the Softmax function is invariant to a common shift in its inputs, we only need to learn the difference in logits. We can define a single logit $\alpha = \alpha_1 - \alpha_2$ and simplify (5) as

$$\begin{aligned} z_1 &= \frac{\exp((\alpha_1 + g_1)/\tau)}{\exp((\alpha_1 + g_1)/\tau) + \exp((\alpha_2 + g_2)/\tau)} \\ &= \frac{1}{1 + \exp\left(\frac{(\alpha_2 + g_2) - (\alpha_1 + g_1)}{\tau}\right)} \\ &= \text{Sigmoid}\left(\frac{\alpha + g_1 - g_2}{\tau}\right). \end{aligned} \quad (6)$$

The simplified expression in (6) is the Gumbel-Sigmoid distribution, which serves as an efficient reparameterization

for binary random variables. We now extend this principle to construct the entire differentiable proxy for the selection matrix Ω . Specifically, each element $[\Omega]_{i,j} \in \{0, 1\}$ of the matrix, representing the connection between the i -th antenna and the j -th RF chain, is modeled as an independent sample from this Gumbel-Sigmoid distribution. Consequently, the trainable parameters of the selection module constitute a logit matrix $\alpha \in \mathbb{C}^{N_{\text{TX}} \times N_{\text{RF}}}$, where each logit $[\alpha]_{i,j}$ individually parameterizes the generation of its corresponding “soft” selection value $[\Omega]_{i,j}$. Therefore, the optimization problem (4) is transformed to

$$\underset{\mathbf{F}_{\text{RF}}, \mathbf{F}_{\text{BB}}, \alpha, \omega}{\text{maximize}} \quad \sum_{k=1}^K R_k + \rho \cdot \text{CRB}^{-1}(\theta_k) \quad (7a)$$

$$\text{subject to} \quad |[\mathbf{F}_{\text{RF}}]_{i,j}| = 1, \forall i, j, \quad (7b)$$

$$\|\tilde{\mathbf{F}}_{\text{RF}} \mathbf{F}_{\text{BB}}\|_F^2 \leq P_t, \quad (7c)$$

$$\omega \in \{0, 1\}^{N_{\text{TX}} \times 1}, \quad (7d)$$

$$\sum_{i=1}^{N_{\text{TX}}} \omega_i = N_{\text{RF}}. \quad (7e)$$

The selection at the receive array imposes a strict cardinality constraint (7e) that exactly N_{RF} out of N_{TX} antenna elements must be activated. To make this combinatorial selection process compatible with gradient-based optimization, we employ a deterministic relaxation technique for “Top- N_{RF} ” selection that directly embeds the cardinality constraint.

Our starting point is the unconstrained logit vector $\eta \in \mathbb{C}^{N_{\text{TX}}}$. To enforce the constraint while controlling the sharpness of the selection, we introduce two components, that is, an adaptive offset λ and a temperature parameter τ . The n -th element of the final “soft” selection vector ω is formulated as

$$\omega_n = \sigma\left(\frac{\eta_n - \lambda(\eta, \tau)}{\tau}\right), \quad (8)$$

where $\sigma(\cdot)$ is the Sigmoid function and $\lambda(\eta, \tau)$ represents that the adaptive offset λ is determined by η and τ . The value of λ is determined by solving the root-finding problem, which also depends on τ and is given by

$$\sum_{n=1}^{N_{\text{TX}}} \sigma\left(\frac{\eta_n - \lambda(\eta, \tau)}{\tau}\right) - N_{\text{RF}} = 0. \quad (9)$$

After that, the optimization problem (7) is transformed to

$$\underset{\mathbf{F}_{\text{RF}}, \mathbf{F}_{\text{BB}}, \alpha, \eta, \lambda}{\text{maximize}} \quad \sum_{k=1}^K R_k + \rho \cdot \text{CRB}^{-1}(\theta_k) \quad (10a)$$

$$\text{subject to} \quad |[\mathbf{F}_{\text{RF}}]_{i,j}| = 1, \forall i, j, \quad (10b)$$

$$\|\tilde{\mathbf{F}}_{\text{RF}} \mathbf{F}_{\text{BB}}\|_F^2 \leq P_t, \quad (10c)$$

$$\sum_{n=1}^{N_{\text{TX}}} \sigma \left(\frac{\eta_n - \lambda}{\tau} \right) = N_{\text{RF}}. \quad (10d)$$

For any given $\boldsymbol{\eta}$ and τ , the left-hand side of (9) remains a monotonically decreasing function of λ . This ensures that a unique solution for λ can be efficiently found by using numerical methods like the bisection algorithm during the forward pass.

A challenge arises in the backward pass, as the root-finding process for λ is not directly differentiable. We leverage the Implicit Function Theorem [36] to compute the gradient. Define $F(\boldsymbol{\eta}, \lambda) = \sum_{n=1}^{N_{\text{TX}}} \sigma \left(\frac{\eta_n - \lambda(\boldsymbol{\eta}, \tau)}{\tau} \right) - N_{\text{RF}}$. Taking the partial derivative of (9) with respect to η_i , according to the chain rule

$$\frac{\partial F}{\partial \eta_i} + \frac{\partial F}{\partial \lambda} \frac{\partial \lambda}{\partial \eta_i} = 0, \quad (11)$$

we can derive

$$\sigma' \left(\frac{\eta_i - \lambda}{\tau} \right) - \sum_{n=1}^{N_{\text{TX}}} \sigma' \left(\frac{\eta_n - \lambda}{\tau} \right) \frac{\partial \lambda}{\partial \eta_i} = 0, \quad (12)$$

where $\sigma'(x) = \sigma(x)(1 - \sigma(x))$ is the derivative function of $\sigma(x)$.

The analytical gradient for λ is given by

$$\frac{\partial \lambda}{\partial \eta_i} = \frac{\sigma' \left(\frac{\eta_i - \lambda}{\tau} \right)}{\sum_{n=1}^{N_{\text{TX}}} \sigma' \left(\frac{\eta_n - \lambda}{\tau} \right)}. \quad (13)$$

This gradient allows error signals to backpropagate through the numerically-solved λ to the network's logits $\boldsymbol{\eta}$.

To facilitate the training process, the temperature τ is gradually annealed from a high initial value to a small final value. This encourages exploration in the early stages and leads to near-binary selections near the end of training. We adopt an exponential decay schedule as

$$\tau^{(t)} = \tau_0 \exp(-ct), \quad (14)$$

where t is the training step, τ_0 is the initial temperature and $c > 0$ is the decay rate.

It should be noted that while the Gumbel-Softmax reparameterization has been widely used for discrete optimization, the relaxation design in this work differs from standard Gumbel-Softmax in two important aspects. First, at the transmitter, instead of applying a two-dimensional categorical Gumbel-Softmax for each antenna–RF-chain connection, we simplify the binary selection into a scalar Gumbel-Sigmoid parameterized by a single logit difference. This reduces the number of trainable parameters and provides a more compact differentiable representation that is convenient for joint optimization with the complex analog beamformer. Second, at the receiver, we design a deterministic Top- N_{RF} relaxation with an

adaptive offset λ . Unlike standard Gumbel-Softmax, which does not naturally enforce a fixed cardinality, the proposed relaxation exactly satisfies the cardinality constraint through the root-finding condition in (9), while the implicit function theorem is used to enable gradient backpropagation through the adaptive offset. These two modifications make the relaxation better suited to the proposed antenna selection problem than direct application of off-the-shelf Gumbel-Softmax.

4 The Proposed AS-Aided HBF Framework

In the preceding section, we address the combinatorial nature of the AS matrices by reparameterizing them as differentiable functions of learnable logits $\boldsymbol{\alpha}$ and $\boldsymbol{\eta}$. While this crucial step makes the selection process amenable to gradient-based methods, the overall optimization problem formulated in (10) remains a significant challenge. The difficulty lies in the joint optimization of the newly introduced selection logits $\boldsymbol{\alpha}$, $\boldsymbol{\eta}$, the analog beamformer $\mathbf{F}_{\text{RF}}$ and the digital beamformer $\mathbf{F}_{\text{BB}}$. These variables are strongly coupled within the sum-rate objective function. To tackle this complex joint problem, we propose an alternating optimization (AO) framework, which decomposes it into a series of more manageable subproblems.

However, the original objective function (10a) is not directly amenable to such a decomposition. To render the problem tractable for the AO framework, we first perform a series of transformations to decouple the variables. First, we apply the Lagrangian dual transform, a technique central to weighted sum-rate (WSR) maximization. By introducing an auxiliary variable $\boldsymbol{\mu}$, this step effectively linearizes the objective with respect to each user's rate, breaking the inter-user interference coupling inherent in the logarithmic structure. This yields the equivalent objective function as

$$\begin{aligned} \mathcal{L} = \sum_{k=1}^K \left\{ \ln(1 + \mu_k) - \mu_k + \frac{(1 + \mu_k) \left| \mathbf{h}_k^H \tilde{\mathbf{F}}_{\text{RF}} \mathbf{v}_k \right|^2}{\sum_{l=1}^K \left| \mathbf{h}_k^H \tilde{\mathbf{F}}_{\text{RF}} \mathbf{v}_l \right|^2 + \sigma_k^2} \right. \\ \left. + \rho \cdot \text{CRB}^{-1}(\theta_k) \right\}. \end{aligned} \quad (15)$$

While the user variables are now decoupled, the resulting expression in (15) still contains a sum of complex multi-ratio terms. At the second step, we apply the fractional programming (FP) transform [37, 38]. This introduces another auxiliary variable $\boldsymbol{\mu}$, and can convert each fractional term into a more tractable quadratic form, resulting in the final

objective function as

$$\mathcal{L} = \sum_{k=1}^K \left\{ \ln(1 + \mu_k) - \mu_k + 2\sqrt{1 + \mu_k} \operatorname{Re} \left\{ \xi_k \mathbf{h}_k^H \tilde{\mathbf{F}}_{\text{RF}} \mathbf{v}_k \right\} - |\xi_k|^2 \left(\sum_{l=1}^K \left| \mathbf{h}_k^H \tilde{\mathbf{F}}_{\text{RF}} \mathbf{v}_l \right|^2 + \sigma_k^2 \right) + \rho \cdot \text{CRB}^{-1}(\theta_k) \right\}. \quad (16)$$

With the objective transformed into the form of (16), the problem is the favorable structure for the AO algorithm. We can iteratively solve for each block of variables including $\mu, \xi, \mathbf{F}_{\text{RF}}, \mathbf{F}_{\text{BB}}, \alpha$ and η until the solution converges.

4.1 Optimization of Transmit and Receive Antenna Selection

4.1.1 Optimization of Transmit Binary Selection Matrix

For fixed values of the other variables, we now address the subproblem of optimizing the transmit selection matrix Ω . As established in Section 3, we optimize a matrix of trainable logits, $\alpha \in \mathbb{C}^{N_{\text{TX}} \times N_{\text{RF}}}$, which generates a differentiable proxy for Ω via the Gumbel-Sigmoid function from (6) as

$$[\Omega]_{i,j} = \sigma \left(\frac{[\alpha]_{i,j} + g_1 - g_2}{\tau} \right), \quad (17)$$

where g_1 and g_2 are i.i.d. samples from a Gumbel(0,1) distribution.

To update the logits α , the gradient of the objective function $\mathcal{L}$ from (16) with respect to α is given by

$$\nabla_{\alpha} \mathcal{L} = \frac{\partial \mathcal{L}}{\partial \Omega} \odot \frac{\partial \Omega}{\partial \alpha}, \quad (18)$$

where the Jacobian term is $\frac{\partial \Omega}{\partial \alpha} = \frac{1}{\tau} \Omega \odot (\mathbf{1}_{N_{\text{TX}} \times N_{\text{RF}}} - \Omega)$. The complete gradient is thus expressed as

$$\begin{aligned} \nabla_{\alpha} \mathcal{L} = & \sum_{k=1}^K \left[2 \left(\sqrt{1 + \mu_k} \xi_k \mathbf{h}_k^H \mathbf{v}_k - |\xi_k|^2 \sum_{l=1}^K \mathbf{h}_k^H \tilde{\mathbf{F}}_{\text{RF}} \mathbf{v}_l \mathbf{h}_k^H \mathbf{v}_l^H \right) \right. \\ & + 4\rho \frac{|\beta_k|^2}{\sigma_s^2} \left(\dot{\mathbf{b}}^H(\theta_k) \dot{\mathbf{b}}(\theta_k) \sum_{l=1}^K \mathbf{a}^H(\theta_k) \tilde{\mathbf{F}}_{\text{RF}} \mathbf{v}_l \mathbf{a}(\theta_k) \mathbf{v}_l^H \right. \\ & \left. \left. + \mathbf{b}^H(\theta_k) \mathbf{b}(\theta_k) \sum_{l=1}^K \dot{\mathbf{a}}^H(\theta_k) \tilde{\mathbf{F}}_{\text{RF}} \mathbf{v}_l \dot{\mathbf{a}}(\theta_k) \mathbf{v}_l^H \right) \right] \\ & \odot \frac{1}{\tau} \Omega \odot (\mathbf{1}_{N_{\text{TX}} \times N_{\text{RF}}} - \Omega). \end{aligned} \quad (19)$$

Remark 3: The gradient expression in (19) reveals two key insights. First, the term $\Omega \odot (\mathbf{1}_{N_{\text{TX}} \times N_{\text{RF}}} - \Omega)$ acts as a natural gating mechanism: when the soft selection Ω approaches 0 or 1, the gradient vanishes, locking in the decision and preventing oscillations. Second, the temperature τ controls the exploration-exploitation trade-off: a high τ suppresses gradients and keeps Ω near 0.5 for exploration, while a low τ amplifies gradients and drives Ω toward binary

values for exploitation. Together with the exponential decay schedule introduced in (14), the proposed method ensures a smooth transition from exploration to exploitation.

To maximize $\mathcal{L}$, we update α by using one step of the gradient ascent method as

$$\alpha^{\text{new}} = \alpha + \kappa_{\alpha} \cdot \nabla_{\alpha} \mathcal{L}, \quad (20)$$

where κ_{α} is the step size. A backtracking line search is employed to adapt κ_{α} , ensuring that the ascent condition $\mathcal{L}(\alpha^{\text{new}}) > \mathcal{L}(\alpha)$ is met. Following the update of α , a new soft matrix Ω is generated by using fresh Gumbel samples for the next iteration of the AO algorithm.

This adaptive update of the transmit selection logits α enables the system to dynamically reconfigure the antenna subset in response to the instantaneous channel conditions, directly addressing the first limitation—the inability of fixed subarrays to adapt to task dynamics.

4.1.2 Optimization of Receive Binary Selection Vector

The optimization of the receive selection vector ω is constrained by a fixed cardinality, requiring the specialized differentiable relaxation introduced in Section 3. The goal is to update the underlying logits $\eta \in \mathbb{C}^{N_{\text{TX}}}$ by maximizing the objective $\mathcal{L}$. Therefore, we use the gradient ascent method, and compute the gradient $\nabla_{\eta} \mathcal{L}$.

The gradient is derived by applying the chain rule as

$$\nabla_{\eta} \mathcal{L} = \left(\frac{\partial \omega}{\partial \eta} \right)^T \frac{\partial \mathcal{L}}{\partial \omega}. \quad (21)$$

This decomposition involves two key components, that is, the Jacobian matrix $\frac{\partial \omega}{\partial \eta}$ and the gradient with respect to ω , $\frac{\partial \mathcal{L}}{\partial \omega}$.

The Jacobian matrix $\frac{\partial \omega}{\partial \eta}$ quantifies how a change in the logits affects the soft selection vector. Using the Implicit Function Theorem given in Section 3, its (i, j) -th element is derived as

$$\frac{\partial \omega_i}{\partial \eta_j} = \frac{1}{\tau} \sigma' \left(\frac{\eta_i + \lambda}{\tau} \right) \left(\delta_{ij} - \frac{\sigma' \left(\frac{\eta_i + \lambda}{\tau} \right)}{\sum_{k=1}^{N_{\text{TX}}} \sigma' \left(\frac{\eta_k + \lambda}{\tau} \right)} \right), \quad (22)$$

where $\delta_{ij} = \{1 \text{ if } i = j, 0 \text{ if } i \neq j\}$.

The partial derivative $\frac{\partial \mathcal{L}}{\partial \omega}$ is the standard Euclidean gradient of the objective function (16) with respect to the soft selection vector ω , and derived as

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial \omega} = & 4\rho \sum_{k=1}^K \frac{|\beta_k|^2}{\sigma_s^2} \left[\pi^2 \cos^2(\theta_k) \tilde{\omega} \sum_{l=1}^K \left| \mathbf{a}^H(\theta_k) \tilde{\mathbf{F}}_{\text{RF}} \mathbf{v}_l \right|^2 \right. \\ & \left. + \sum_{l=1}^K \omega \left| \dot{\mathbf{a}}^H(\theta_k) \tilde{\mathbf{F}}_{\text{RF}} \mathbf{v}_l \right|^2 \right], \end{aligned} \quad (23)$$

where $\tilde{\omega} = [m_1, m_2, \dots, m_{N_{\text{TX}}}]^T \odot \omega$ and m_i is the position of the i -th element normalized by the inter-element spacing

d.

By computing these two components and combining them as shown in (21), we can obtain the gradient $\nabla_{\eta}\mathcal{L}$. The logits are then updated using a gradient ascent step as

$$\boldsymbol{\eta}^{\text{new}} = \boldsymbol{\eta} + \kappa_{\eta} \cdot \nabla_{\eta}\mathcal{L}, \quad (24)$$

where κ_{η} is a step size determined by a backtracking line search to ensure an increase of the objective function. The updated logits $\boldsymbol{\eta}^{\text{new}}$ will then be used to generate the new selection vector $\boldsymbol{\omega}$, which is forward to the next AO iteration.

By enabling gradient-based optimization under the strict cardinality constraint, this module provides the necessary hardware elasticity for the receive array, allowing the system to dynamically activate the most informative antenna elements while maintaining compatibility with the limited number of RF chains.

4.2 Optimization of Other Variables

4.2.1 Optimization of Auxiliary Variables

With the beamformers and selection schemes fixed, the auxiliary variables $\boldsymbol{\mu}$ and $\boldsymbol{\xi}$ are updated via their optimal closed-form solutions [37], derived by setting the first-order derivatives of (16) to zero as

$$\mu_k^* = \frac{\left| \mathbf{h}_k^H \tilde{\mathbf{F}}_{\text{RF}} \mathbf{v}_k \right|^2}{\sum_{l \neq k} \left| \mathbf{h}_k^H \tilde{\mathbf{F}}_{\text{RF}} \mathbf{v}_l \right|^2 + \sigma_k^2}, \quad (25)$$

$$\xi_k^* = \frac{\sqrt{1 + \mu_k} \mathbf{h}_k^H \tilde{\mathbf{F}}_{\text{RF}} \mathbf{v}_k}{\sum_{l=1}^K \left| \mathbf{h}_k^H \tilde{\mathbf{F}}_{\text{RF}} \mathbf{v}_l \right|^2 + \sigma_k^2}. \quad (26)$$

4.2.2 Optimization of Digital Beamformer

With other variables held fixed, the optimization subproblem for the digital beamformer $\mathbf{F}_{\text{BB}}$ is formulated under the power constraint (10c). The Lagrangian function is constructed as

$$\tilde{\mathcal{L}}(\mathbf{F}_{\text{BB}}, \lambda) = \mathcal{L} - \lambda (\text{Tr}(\mathbf{F}_{\text{BB}}^H \mathbf{A} \mathbf{F}_{\text{BB}}) - P_t), \quad (27)$$

where $\mathbf{A} = \tilde{\mathbf{F}}_{\text{RF}}^H \tilde{\mathbf{F}}_{\text{RF}}$ is defined for the notational simplicity, and $\lambda \geq 0$ is the Lagrange multiplier associated with the power constraint.

The optimal solution to this problem must satisfy the Karush-Kuhn-Tucker (KKT) conditions. The key conditions are stationarity and complementary slackness as

$$\frac{\partial \tilde{\mathcal{L}}(\mathbf{F}_{\text{BB}}, \lambda)}{\partial \mathbf{F}_{\text{BB}}} = \mathbf{0}, \quad (28)$$

$$\lambda (\text{Tr}(\mathbf{F}_{\text{BB}}^H \mathbf{A} \mathbf{F}_{\text{BB}}) - P_t) = 0. \quad (29)$$

From the stationarity condition (28), the k -th column of $\mathbf{F}_{\text{BB}}$ is obtained as

$$\mathbf{v}_k(\lambda) = \sqrt{1 + \mu_k \xi_k^*} \mathbf{A}^{-1}(\lambda) \tilde{\mathbf{F}}_{\text{RF}}^H \mathbf{h}_k, \quad (30)$$

where

$$\begin{aligned} \mathbf{A}(\lambda) = & \sum_{l=1}^K |\xi_l|^2 \tilde{\mathbf{F}}_{\text{RF}}^H \mathbf{h}_l \mathbf{h}_l^H \tilde{\mathbf{F}}_{\text{RF}} \\ & - \rho \sum_{l=1}^K \frac{2|\beta_l|^2}{\sigma_s^2} \left(\dot{\mathbf{b}}^H(\theta_l) \dot{\mathbf{b}}(\theta_l) \tilde{\mathbf{F}}_{\text{RF}}^H \mathbf{a}(\theta_l) \mathbf{a}^H(\theta_l) \tilde{\mathbf{F}}_{\text{RF}} \right. \\ & \left. + \mathbf{b}^H(\theta_l) \mathbf{b}(\theta_l) \tilde{\mathbf{F}}_{\text{RF}}^H \dot{\mathbf{a}}(\theta_l) \dot{\mathbf{a}}^H(\theta_l) \tilde{\mathbf{F}}_{\text{RF}} \right) + \lambda \tilde{\mathbf{F}}_{\text{RF}}^H \tilde{\mathbf{F}}_{\text{RF}}. \end{aligned} \quad (31)$$

The optimal Lagrange multiplier λ^* is determined via the complementary slackness condition (29). Two cases are considered: (1) If the solution with $\lambda = 0$ already satisfies the power constraint, $\lambda^* = 0$; (2) Otherwise, $\lambda^* > 0$ is derived by solving $g(\lambda) = \text{Tr}(\mathbf{F}_{\text{BB}}^H(\lambda) \mathbf{A} \mathbf{F}_{\text{BB}}(\lambda)) - P_t = 0$ via bisection search. The optimal beamformer $\mathbf{F}_{\text{BB}}^*$ is then obtained by substituting λ^* into (30).

This closed-form digital beamformer update leverages the high-precision angle estimates from the sensing loop to maximize the weighted sum-rate, thereby closing the sensing-aided communication loop and improving the overall beam quality.

4.2.3 Optimization of Analog Beamformer

For fixed values of the other variables, the subproblem for optimizing the analog beamformer $\mathbf{F}_{\text{RF}}$ is constrained by the non-convex unit-modulus constraint (10b). We address this using manifold optimization, recognizing that the feasible set forms a complex circle manifold $\mathcal{M}^{N_{\text{TX}} \times N_{\text{RF}}}$ with $\mathcal{M} = \{z \in \mathbb{C} : |z| = 1\}$. The update proceeds as follows.

Step 1: Computing the Euclidean Gradient. First, we compute the standard (Euclidean) gradient of the objective function $\mathcal{L}$ from (16) with respect to $\mathbf{F}_{\text{RF}}$. The gradient, denoted as $\nabla_{\mathbf{F}_{\text{RF}}} \mathcal{L}$, is given by

$$\begin{aligned} \nabla_{\mathbf{F}_{\text{RF}}} \mathcal{L} = & 2 \sum_{k=1}^K \left[\sqrt{1 + \mu_k} \xi_k \mathbf{h}_k \mathbf{v}_k^H \right. \\ & \left. - |\xi_k|^2 \sum_{l=1}^K \mathbf{h}_k^H \tilde{\mathbf{F}}_{\text{RF}} \mathbf{v}_l \mathbf{h}_k \mathbf{v}_l^H \right] \\ & + 4\rho \sum_{k=1}^K \frac{|\beta_k|^2}{\sigma_s^2} \left[\dot{\mathbf{b}}^H(\theta_k) \dot{\mathbf{b}}(\theta_k) \sum_{l=1}^K \mathbf{a}^H(\theta_k) \tilde{\mathbf{F}}_{\text{RF}} \mathbf{v}_l \mathbf{a}(\theta_k) \mathbf{v}_l^H \right. \\ & \left. + \mathbf{b}^H(\theta_k) \mathbf{b}(\theta_k) \sum_{l=1}^K \dot{\mathbf{a}}^H(\theta_k) \tilde{\mathbf{F}}_{\text{RF}} \mathbf{v}_l \dot{\mathbf{a}}(\theta_k) \mathbf{v}_l^H \right]. \end{aligned} \quad (32)$$

Step 2: Projecting to the Tangent Space. A standard gradient update step, $\mathbf{F}_{\text{RF}} - \alpha \cdot \nabla_{\mathbf{F}_{\text{RF}}} \mathcal{L}$, would move the point off the manifold. To maintain the unit-modulus constraint, we project the Euclidean gradient onto the tangent space of the manifold at the current point $\mathbf{F}_{\text{RF}}$. For the product manifold $\mathcal{M}^{N_{\text{TX}} \times N_{\text{RF}}}$, this projection is performed element-wise. The

resulting Riemannian gradient, $\text{grad}_{\mathbf{F}_{\text{RF}}} \mathcal{L}$ is computed as

$$\text{grad}_{\mathbf{F}_{\text{RF}}} \mathcal{L} = \text{Proj}_{\mathbf{F}_{\text{RF}}} \nabla \mathcal{L} = \nabla \mathcal{L} - \text{Re} \{ \nabla \mathcal{L} \odot \mathbf{F}_{\text{RF}} \} \odot \mathbf{F}_{\text{RF}}. \quad (33)$$

This ensures the search direction lies within the tangent space, locally respecting the manifold's curvature.

Step 3: Retraction to the Manifold. After determining the search direction with the Riemannian gradient, we update the current point by moving along this direction in the tangent space with a step size of α . The resulting point, $\mathbf{F}'_{\text{RF}} = \mathbf{F}_{\text{RF}} - \alpha \cdot \text{grad}_{\mathbf{F}_{\text{RF}}} \mathcal{L}$, may lie slightly outside the manifold. We map it back onto $\mathcal{M}^{N_{\text{TX}} \times N_{\text{RF}}}$ using a retraction operation. For the complex circle manifold, a common retraction is the element-wise projection as

$$[\mathbf{F}_{\text{RF}}]_{ij}^{\text{new}} = \frac{[\mathbf{F}'_{\text{RF}}]_{ij}}{[\mathbf{F}'_{\text{RF}}]_{ij}} = \frac{[\mathbf{F}_{\text{RF}} - \alpha \cdot \text{grad} \mathcal{L}]_{ij}}{[\mathbf{F}_{\text{RF}} - \alpha \cdot \text{grad} \mathcal{L}]_{ij}}. \quad (34)$$

The step size α is determined using a standard backtracking line search method, such as the Armijo rule, to ensure sufficient decrease of the objective function on the manifold.

The manifold optimization ensures that the analog beamformer satisfies the unit-modulus constraint while aligning the beam directions with the estimated UAV angles, further improving the beam quality and interference suppression in dynamic environments.

The overall joint transceiver optimization problem is solved using the AO framework, which iteratively updates each block of variables while keeping the others fixed. The complete procedure, encapsulating the solutions detailed in Section 4, is summarized in Algorithm 1. The gradient-based updates for the selection logits (α and η) and the analog beamformer $\mathbf{F}_{\text{RF}}$ implicitly include a backtracking line search to determine an appropriate step size (κ_α , κ_η and α) that guarantees an increase in the objective function value in each iteration.

4.3 Convergence and Complexity Analysis

A theoretical analysis of the proposed algorithm is provided. We first establish its convergence to a stationary point of the objective function and then analyze its computational complexity.

4.3.1 Convergence Analysis

Let $\Theta^{(t)} = \{\mu^{(t)}, \xi^{(t)}, \mathbf{F}_{\text{RF}}^{(t)}, \mathbf{F}_{\text{BB}}^{(t)}, \alpha^{(t)}, \eta^{(t)}\}$ be the set of all variables at iteration t , and let $\mathcal{L}(\Theta^{(t)}; \tau^{(t)})$ be the value of the objective function. The proposed algorithm updates both the variables Θ and the temperature τ at each iteration. To guarantee the convergence, we first derive the following properties.

Algorithm 1 Proposed Alternating Optimization for AS-Aided HBF

- 1: **Input:** $K, N_{\text{TX}}, N_{\text{RF}}, P_t, \{\theta_k\}_{k=1}^K$, weighting factor ρ , max iterations T_{max} , convergence threshold ϵ , temperature parameters $\tau_0, \tau_{\text{min}}$, and annealing rate c .
- 2: **Output:** $\mathbf{F}_{\text{RF}}, \mathbf{F}_{\text{BB}}, \Omega$, and ω .
- 3: **Initialize:**
- 4: Set iteration counter $t = 0$.
- 5: Initialize $\tau = \tau_0$.
- 6: Initialize randomly $\mathbf{F}_{\text{RF}}^{(0)}, \mathbf{F}_{\text{BB}}^{(0)}, \alpha^{(0)}$, and $\eta^{(0)}$.
- 7: Generate $\Omega^{(0)}$ via (17) and $\omega^{(0)}$ via (8).
- 8: Calculate initial objective value $\mathcal{L}^{(0)}$ using (16).
- 9: **repeat**
- 10: $t \leftarrow t + 1$
- 11: Compute gradient $\nabla_\alpha \mathcal{L}$ at $\alpha^{(t-1)}$ using (19).
- 12: Update logits via gradient ascent: $\alpha^{(t)} \leftarrow \alpha^{(t-1)} + \kappa_\alpha \cdot \nabla_\alpha \mathcal{L}$.
- 13: Generate new soft selection matrix $\Omega^{(t)}$ from $\alpha^{(t)}$ using (17).
- 14: Compute gradient $\nabla_\eta \mathcal{L}$ at $\eta^{(t-1)}$ using the chain rule in (21).
- 15: Update logits via gradient ascent: $\eta^{(t)} \leftarrow \eta^{(t-1)} + \kappa_\eta \cdot \nabla_\eta \mathcal{L}$.
- 16: Generate new soft selection vector $\omega^{(t)}$ from $\eta^{(t)}$ via (8).
- 17: Update τ via (14).
- 18: **for** $k = 1, \dots, K$ **do**
- 19: Update $\mu_k^{(t)}$ via (25) with variables from iteration $t - 1$.
- 20: Update $\xi_k^{(t)}$ via (26) with variables from iteration $t - 1$.
- 21: **end for**
- 22: Update $\mathbf{F}_{\text{BB}}^{(t)}$ by solving the KKT system based on (30).
- 23: Compute Euclidean gradient $\nabla_{\mathbf{F}_{\text{RF}}} \mathcal{L}$ at $\mathbf{F}_{\text{RF}}^{(t-1)}$ using (32).
- 24: Project gradient $\nabla_{\mathbf{F}_{\text{RF}}} \mathcal{L}$ to tangent space using (33).
- 25: Update $\mathbf{F}_{\text{RF}}^{(t)}$ via retraction with step α using (34).
- 26: Calculate the new objective value $\mathcal{L}^{(t)}$ with all updated variables.
- 27: **until** $\frac{|\mathcal{L}^{(t)} - \mathcal{L}^{(t-1)}|}{|\mathcal{L}^{(t-1)}|} < \epsilon$ **or** $t \geq T_{\text{max}}$
- 28: **Finalization:**
- 29: Obtain final binary Ω by element-wise thresholding of $\Omega^{(t)}$ at 0.5.
- 30: Obtain final binary ω by selecting indices of the N_{RF} largest elements in $\eta^{(t)}$.

Property 1 (Boundedness). For any $\tau \geq 0$, the objective function $\mathcal{L}(\Theta; \tau)$ is uniformly bounded from above over the feasible set of Θ . That is, there exists a supremum L_{sup} such

that $\mathcal{L}(\Theta; \tau) \leq L_{\text{sup}} < \infty$. This is physically justified as transmit power is finite, leading to finite channel capacity and CRB.

Property 2 (Smoothness). *The function $\mathcal{L}(\Theta; \tau)$ is continuously differentiable with respect to both Θ and τ . Furthermore, $\mathcal{L}$ is Lipschitz continuous with respect to τ , i.e., there exists a constant $K_\tau > 0$ such that for any feasible Θ and any $\tau_1, \tau_2 \geq 0$:*

$$|\mathcal{L}(\Theta; \tau_1) - \mathcal{L}(\Theta; \tau_2)| \leq K_\tau |\tau_1 - \tau_2|. \quad (35)$$

This ensures that small changes in temperature do not cause unbounded changes in the objective function.

Property 3 (Annealing Schedule). *The temperature sequence $\{\tau^{(t)}\}$ is non-negative and satisfies: (a) $\lim_{t \rightarrow \infty} \tau^{(t)} = 0$. (b) The total variation is finite: $\sum_{t=0}^{\infty} |(\tau^{(t+1)} - \tau^{(t)})| < \infty$. This condition enforces that the temperature anneals sufficiently slowly. The decay schedule proposed in (14) satisfies this condition.*

With these properties, the convergence of the proposed algorithm can be stated as follows:

Theorem 1. *Under properties 1-3, the sequence of variables generated by the proposed AO algorithm is guaranteed to converge to a stationary point of the objective function.*

Proof. See Appendix B. $\square$

It should be noted that the convergence result in Theorem 1 applies to the differentiable relaxed problem (10) rather than the original mixed-integer problem (4). Since the Gumbel-Softmax relaxation provides a continuous surrogate of the binary selection variables, the final AS is obtained by projecting the relaxed solution onto the feasible binary set after the temperature annealing process. As $\tau \rightarrow 0$, the Gumbel-Sigmoid output converges to a Bernoulli sample, so that the relaxed variables approach exact binary values. Consequently, the projection loss becomes negligible. A quantitative evaluation of this loss is provided in Section 6.6.

4.3.2 Complexity Analysis

The complexity analysis of the proposed AO algorithm per iteration considers the dominant operations in each variable update step, with particular attention to the problem dimensions, that is, N_{TX} (number of transmit antennas), N_{RF} (number of RF chains), and K (number of UAVs).

(1) Transmit Antenna Selection Ω : The gradient computation in (19) involves matrix multiplications and Hadamard products with complexity $\mathcal{O}(KN_{\text{TX}}N_{\text{RF}})$.

Table 1 Complexity Comparison

Algorithms	Complexity
Proposed AO Algorithm	$\mathcal{O}(N_{\text{TX}}^2 + KN_{\text{TX}}N_{\text{RF}}^2 + LN_{\text{RF}}^3)$
Exhaustive Search	$\mathcal{O}\left(2^{N_{\text{TX}}N_{\text{RF}}}\binom{N_{\text{TX}}}{N_{\text{RF}}}N_{\text{TX}}\right)$
Genetic Algorithm [16]	$\mathcal{O}(I_{\text{GA}}PN_{\text{TX}}^2)$
Transmitter-Only Optimization	$\mathcal{O}(KN_{\text{TX}}N_{\text{RF}}^2)$

(2) Receive Antenna Selection ω : The Jacobian matrix computation in (22) dominates this step with complexity $\mathcal{O}(N_{\text{TX}}^2)$, while the gradient $\partial\mathcal{L}/\partial\omega$ in (23) involves $\mathcal{O}(KN_{\text{TX}})$.

(3) Auxiliary Variables μ, ξ : The closed-form updates in (25) and (26) require $\mathcal{O}(KN_{\text{TX}}N_{\text{RF}})$ operations.

(4) Digital Beamformer $\mathbf{F}_{\text{BB}}$: Constructing $\mathbf{A}(\lambda)$ in (31) requires $\mathcal{O}(KN_{\text{TX}}N_{\text{RF}}^2)$ operations. Solving the linear system and the bisection search for λ^* require $\mathcal{O}(LN_{\text{RF}}^3)$, where L is the number of bisection iterations.

(5) Analog Beamformer $\mathbf{F}_{\text{RF}}$: The Euclidean gradient (32), tangent space projection (33), and retraction (34) collectively require $\mathcal{O}(KN_{\text{TX}}N_{\text{RF}})$ operations.

For typical massive MIMO configurations where $N_{\text{TX}} \gg N_{\text{RF}}$, the dominant term is $\mathcal{O}(N_{\text{TX}}^2)$ due to the receive selection Jacobian, followed by $\mathcal{O}(KN_{\text{TX}}N_{\text{RF}}^2)$ from the digital beamformer design. Therefore, the overall per-iteration complexity is $\mathcal{O}(N_{\text{TX}}^2 + KN_{\text{TX}}N_{\text{RF}}^2 + LN_{\text{RF}}^3)$.

To provide a clearer perspective on the algorithm's efficiency, we compare its complexity with three benchmarks, including Exhaustive Search, Genetic Algorithm (with I_{GA} and P denoting the number of iterations and population size, respectively) [16] and Transmitter-Only Optimization. The comparison is shown in Table 1.

The proposed algorithm thus achieves a favorable trade-off, that is, it transforms an intractable combinatorial problem into a tractable optimization with polynomial complexity, while maintaining near-optimal performance through joint transceiver design. Compared to heuristic methods like GA, the proposed gradient-based approach provides theoretical convergence guarantees and more efficient exploration of the solution space. This represents a significant advancement over conventional approaches that either suffer from exponential complexity, rely on non-guaranteed heuristics, or achieve only partial optimization.

5 Angle Estimation

To close the sensing loop by providing accurate angular information for the next transmission cycle in the proposed dynamic AS framework, a robust angle estimation is needed at the receiver side. This task is challenging due to three

factors inherent in our dynamic AS framework, that is, the non-uniform receive array invalidates standard algorithms, the potential coherence of reflected signals degrades subspace methods like the Multiple Signal Classification (MUSIC) [39], and the single-snapshot requirement for tracking rules out multi-snapshot covariance-based estimation. To overcome these joint constraints, we employ the iterative adaptive approach (IAA). Its suitability is threefold: as a single-snapshot method, it meets the real-time tracking requirement; it is independent of array geometry, making it robust to our non-uniform array; and it inherently handles coherent sources by iteratively building a model-based covariance matrix for effective decorrelation.

The core idea is to reframe the angle estimation task as a sparse signal recovery problem. We discretize the entire angular surveillance space of interest, for instance $[-90^\circ, 90^\circ]$, into a fine grid of N_g discrete points, where $N_g \gg K$. Let this grid be $\boldsymbol{\theta} = [\theta_1, \dots, \theta_{N_g}]$.

The received signal vector $\mathbf{r} \in \mathbb{C}^{N_{\text{Tx}} \times 1}$ at the BS, as expressed in (2), can then be modeled as a linear combination of steering vectors corresponding to all potential grid points as

$$\mathbf{r} = \sum_n^{N_g} \mathbf{b}(\theta_n) \tilde{s}_n + \tilde{\mathbf{n}} = \sum_n^{N_g} \mathbf{B}(\boldsymbol{\theta}) \tilde{\mathbf{s}} + \tilde{\mathbf{n}}, \quad (36)$$

where $\tilde{\mathbf{s}} \in \mathbb{C}^{N_g \times 1}$ is a sparse vector representing the complex amplitudes of the signals from each grid point; the true UAV angles correspond to the few non-zero entries in $\tilde{\mathbf{s}}$, and the goal is to recover this sparse vector; $\mathbf{b}(\theta_n) = \boldsymbol{\omega} \odot \mathbf{a}(\theta_n) \in \mathbb{C}^{N_g \times 1}$ is the effective steering vector for the selected receive array at the grid angle θ_n ; $\mathbf{B}(\boldsymbol{\theta}) = [\mathbf{b}(\theta_1), \dots, \mathbf{b}(\theta_{N_g})] \in \mathbb{C}^{N_{\text{Tx}} \times N_g}$ is the dictionary or sensing matrix, constructed from the effective steering vectors for all grid points; $\tilde{\mathbf{n}}$ is the sampled noise vector from (2). Our objective is to estimate the power spectral density across the grid, i.e., $P_n = E[|\tilde{s}_n|^2]$ for $n = 1, \dots, N_g$, from which the UAV angles can be identified as the locations of the K largest peaks.

The fundamental idea of IAA is to estimate the signal amplitude $\tilde{s}_n$ from a particular direction θ_n by treating all other signals ($\tilde{s}_j$ where $j \neq n$) as structured interference. In each iteration, it uses the power estimates from the previous step to build a better estimate of this interference-plus-noise covariance matrix, thereby improving the signal estimation in the current step.

The model in (36) can be rewritten as

$$\mathbf{r} = \mathbf{b}(\theta_n) \tilde{s}_n + \sum_{j \neq n} \mathbf{b}(\theta_j) \tilde{s}_j + \tilde{\mathbf{n}}. \quad (37)$$

The second term represents the combined interference from

all other grid points. Assuming the signals $\tilde{s}_j$ from different grid points are uncorrelated, and also uncorrelated with the noise, the covariance matrix of the second term plus the thermal noise $\mathbf{R}_n$ can be expressed as

$$\begin{aligned} \mathbf{R}_n &= \sum_{j \neq n} E[|\tilde{s}_j|^2] \mathbf{b}(\theta_j) \mathbf{b}^H(\theta_j) + E[\tilde{\mathbf{n}} \tilde{\mathbf{n}}^H] \\ &= \sum_{j \neq n} P_j \mathbf{b}(\theta_j) \mathbf{b}^H(\theta_j) + \sigma_s^2 \text{diag}(\boldsymbol{\omega}), \end{aligned} \quad (38)$$

We can now estimate $\tilde{s}_n$ using a Capon beamformer. The solution of the weight vector $\mathbf{w}_n$ is

$$\mathbf{w}_n = \frac{\mathbf{R}_n^{-1} \mathbf{b}(\theta_n)}{\mathbf{b}^H(\theta_n) \mathbf{R}_n^{-1} \mathbf{b}(\theta_n)}. \quad (39)$$

This gives the estimate for $\tilde{s}_n$ as

$$\hat{s}_n = \mathbf{w}_n^H \mathbf{r} = \frac{\mathbf{b}^H(\theta_n) \mathbf{R}_n^{-1} \mathbf{r}}{\mathbf{b}^H(\theta_n) \mathbf{R}_n^{-1} \mathbf{b}(\theta_n)}. \quad (40)$$

The problem is that $\mathbf{R}_n$ is not known because it depends on the unknown signal powers $\{P_j\}_{j=1}^{N_g}$. IAA resolves this with an iterative procedure. Defining $P_j^{(i)}$ as the power estimate of the j -th signal at iteration i , we can construct an estimate of the total covariance matrix $\mathbf{R}$ as

$$\begin{aligned} \mathbf{R}^{(i)} &= \sum_{j=1}^{N_g} P_j^{(i-1)} \mathbf{b}(\theta_j) \mathbf{b}^H(\theta_j) + \sigma_s^2 \text{diag}(\boldsymbol{\omega}) \\ &= \mathbf{B}(\boldsymbol{\theta}) \mathbf{P}^{(i-1)} \mathbf{B}^H(\boldsymbol{\theta}) + \sigma_s^2 \text{diag}(\boldsymbol{\omega}), \end{aligned} \quad (41)$$

where $\mathbf{P}^{(i-1)}$ is a diagonal matrix with $\{P_j^{(i-1)}\}_{j=1}^{N_g}$ on its diagonal. Applying the matrix inversion lemma to $\mathbf{R}$, we can express $\mathbf{R}_n^{-1}$ in terms of $\mathbf{R}^{-1}$. Plugging this back into the estimator for $\tilde{s}_n$ and after some algebraic simplification, the IAA update rule can be obtained as

$$\hat{s}_n^{(i)} = \frac{\mathbf{b}^H(\theta_n) (\mathbf{R}^{(i)})^{-1} \mathbf{r}}{\mathbf{b}^H(\theta_n) \mathbf{R}_n^{-1} \mathbf{b}(\theta_n)}. \quad (42)$$

After computing the new complex amplitude estimate $\hat{s}_n^{(i)}$, we update its power for the next iteration as $P_n^{(i)} = |\hat{s}_n^{(i)}|^2$. This process is repeated for all $n = 1, \dots, N_g$ and for several iterations until the power estimates converge. The complete procedure for IAA-based angle estimation is summarized in Algorithm 2.

Upon obtaining the converged spatial power spectrum $\{P_n\}_{n=1}^{N_g}$, the angle estimates $\{\theta_k\}_{k=1}^K$ for the K UAVs are determined by finding the angular locations of the K most dominant peaks in the spectrum. These high-precision angles are then utilized to update the transceiver design for the subsequent transmission, thus closing the loop of the proposed ISAC framework.

Algorithm 2 The Proposed IAA Algorithm for Angle Estimation

- 1: **Input:** Received signal $\mathbf{r} \in \mathbb{C}^{N_{\text{Tx}} \times 1}$, receive antenna selection vector $\boldsymbol{\omega}$, noise variance σ_s^2 , number of grid points N_g , convergence threshold ϵ .
- 2: **Preprocessing:**
- 3: Construct the angular grid $\boldsymbol{\theta} = [\theta_1, \dots, \theta_{N_g}]$.
- 4: Construct the dictionary $\mathbf{B}(\boldsymbol{\theta}) = [\mathbf{b}(\theta_1), \dots, \mathbf{b}(\theta_{N_g})]$.
- 5: **Initialization:**
- 6: Set iteration counter $t \leftarrow 0$.
- 7: Initialize the power spectrum using a matched filter:
- 8: $P_n^{(0)} \leftarrow |\mathbf{b}^H(\theta_n)\mathbf{r}|^2 / \|\mathbf{b}(\theta_n)\|^4$ for $n = 1, \dots, N_g$.
- 9: **repeat**
- 10: $i \leftarrow i + 1$
- 11: Construct the diagonal power matrix:
- 12: $\mathbf{P}^{(i-1)} = \text{diag}(P_1^{(i-1)}, \dots, P_{N_g}^{(i-1)})$.
- 13: Compute the covariance matrix:
- 14: $\mathbf{R}^{(i)} = \mathbf{B}(\boldsymbol{\theta})\mathbf{P}^{(i-1)}\mathbf{B}^H(\boldsymbol{\theta}) + \sigma_s^2 \cdot \text{diag}(\boldsymbol{\omega})$.
- 15: Invert the covariance matrix to get $(\mathbf{R}^{(i)})^{-1}$.
- 16: **for** $n = 1$ to N_g **do**
- 17: Compute the complex amplitude:
- 18: $\hat{s}_n^{(i)} = (\mathbf{b}^H(\theta_n)(\mathbf{R}^{(i)})^{-1}\mathbf{r}) / (\mathbf{b}^H(\theta_n)(\mathbf{R}^{(i)})^{-1}\mathbf{b}(\theta_n))$.
- 19: Update the power estimate: $P_n^{(i)} \leftarrow |\hat{s}_n^{(i)}|^2$.
- 20: **end for**
- 21: **until** $\|\mathbf{P}^{(i)} - \mathbf{P}^{(i-1)}\|_F / \|\mathbf{P}^{(i-1)}\|_F < \epsilon$ or max iterations reached.
- 22: **Output:** Power spectrum $\{P_n\}_{n=1}^{N_g} \leftarrow \{P_n^{(i)}\}_{n=1}^{N_g}$.

6 Simulation Results

6.1 Simulation Setup and Parameter Settings

This section provides a comprehensive performance evaluation of the proposed AS-aided HBF framework for ISAC systems. The BS operates at a carrier frequency of 30 GHz with a system bandwidth of 100 MHz. It serves $K = 3$ UAVs located at angles $[-50^\circ, 10^\circ, 40^\circ]$, with distances d_k uniformly distributed in $[80, 150]$ m. Consistent with the channel model in Section 2, the RCS $\sigma_{\text{rcs},k}$ is modeled as a complex Gaussian variable with unit variance. The noise power is set to $\sigma_s^2 = \sigma_k^2 = -84$ dBm, based on a noise figure of 10 dB. In the simulation, the BS is equipped with $N_{\text{Tx}} = 64$ antennas and $N_{\text{RF}} = 6$ RF chains.

To explicitly assess the contributions of the transmitter-side and receiver-side antenna selection mechanisms, we consider four configurations: 1) the proposed joint transmitter-receiver (“Joint Tx-Rx”) scheme, where both the transmit and receive antenna selections are jointly optimized; 2) “Tx-only”, where

antenna selection is restricted to the transmitter side while the receiver array is fixed, a configuration also considered in prior antenna-selection studies [24]; 3) “Rx-only”, where antenna selection is restricted to the receiver side while the transmitter array is fixed; and 4) “Full-array”, which represents a conventional full-array HBF configuration with all antennas active and no antenna selection, as commonly considered in existing HBF-ISAC systems [15]. For a fair comparison, all four configurations employ the same AO-based HBF optimization framework, with only the antenna-selection variables and their corresponding constraints being modified. This four-way comparison provides an ablation study that isolates the individual contributions of the Tx-side and Rx-side selection mechanisms and evaluates the additional benefit of their joint optimization.

6.2 Performance Evaluation and Ablation Study

Fig. 3 illustrates the aggregate spectral efficiency versus the power budget, which refers to the total transmit power constraint P_t in (4c). The proposed Joint Tx-Rx configuration consistently achieves the highest spectral efficiency among the four configurations, followed by the Tx-only, Rx-only, and Full-array configurations. This performance hierarchy provides an ablation study of the proposed framework, showing the respective contributions of transmitter- and receiver-side antenna selection. Specifically, the improvement of Joint Tx-Rx over Tx-only demonstrates the complementary communication gain brought by receiver-side antenna selection, while the larger gap between Tx-only and Rx-only indicates that transmitter-side antenna selection contributes more significantly to communication performance under the considered system configuration.

The communication gain of the antenna-selection configurations over the Full-array configuration reveals the benefit of adaptive antenna selection under the limited-RF-chain architecture. First, by concentrating transmit power onto an optimized subset, the configurations enhance the effective radiated power and thus the received signal strength. Second, since the system’s spatial degrees of freedom are limited by the RF chains, not by the antennas, the proposed selection framework ensures an optimal mapping between these two resources, thereby maximizing beamforming quality. Consequently, the Joint Tx-Rx configuration achieves the highest spectral efficiency by jointly exploiting the spatial degrees of freedom at both the transmitter and receiver, whereas the Tx-only and Rx-only configurations provide only partial gains by optimizing one side of the transceiver.




Fig. 3 Spectral efficiency versus power budget across four antenna selection configurations.

Fig. 4 validates the effectiveness of the IAA-based angle estimator: the dashed MSE curves closely track the solid CRB curves across the entire power range, confirming near-CRB performance. The four configurations exhibit a clear sensing-performance hierarchy, where the Full-array configuration achieves the lowest CRB, followed by the Joint Tx-Rx, Rx-only, and Tx-only configurations. This result complements the spectral-efficiency observations in Fig. 3 and provides an ablation study of the sensing contributions from transmitter and receiver-side antenna selection. Compared with Tx-only and Rx-only, the lower CRB achieved by the Joint Tx-Rx configuration demonstrates that jointly optimizing the transmit and receive antenna selections can partially compensate for the sensing degradation introduced by antenna selection while maintaining superior communication performance.

Sensing accuracy is intrinsically linked to the physical aperture of the receive array. The full-array provides the maximum aperture and thus the lowest theoretical CRB. In contrast, any antenna selection creates a sparse array, which inherently degrades this sensing limit. Among the antenna-selection configurations, the Joint Tx-Rx configuration consistently achieves a lower CRB than both the Tx-only and Rx-only configurations, indicating that coordinated transmitter-receiver optimization can better balance the communication objective and sensing accuracy than optimizing either side alone. The proposed framework is designed not to minimize CRB in isolation, but to optimize a weighted objective. It strategically allocates limited resources (RF chains) to prioritize communication performance, accepting a calculated penalty in sensing. As will be further demonstrated in Fig. 7, this coordinated optimization enables the Joint Tx-Rx configuration to achieve




Fig. 4 CRB versus power budget across four antenna selection configurations.

the most favorable communication-sensing trade-off among the four configurations.

Fig. 5 demonstrates that the performance of all configurations improves as the number of RF chains (N_{RF}) increases, a trend that highlights the critical role of RF chains in defining a system's spatial degrees of freedom. For communication, a larger N_{RF} enables more precise beamforming and interference suppression, directly boosting spectral efficiency. For sensing, it allows for higher-dimensional signal sampling, which enriches the Fisher information and lowers the CRB.

The relative performance hierarchy observed in Figs. 3 and 4 remains consistent as N_{RF} varies, providing a robustness validation of the proposed ablation study. In terms of spectral efficiency, the Joint Tx-Rx configuration consistently achieves the highest performance, followed by the Tx-only, Rx-only, and Full-array configurations. In terms of sensing accuracy, the Full-array configuration maintains the lowest CRB, while the Joint Tx-Rx configuration consistently achieves a lower CRB than both the Tx-only and Rx-only configurations. These results indicate that increasing the number of RF chains benefits all configurations, but the additional spatial degrees of freedom are utilized most effectively through coordinated transmitter-receiver antenna selection. The superior performance of the Joint Tx-Rx configuration demonstrates that jointly optimizing the transmit and receive antenna selections provides a more effective utilization of each additional RF chain for balancing communication and sensing performance.

Fig. 6 dissects the trade-off characteristics of the four configurations by analyzing their performance against the trade-off coefficient, ρ , revealing profound differences in their controllability and efficiency. The proposed Joint Tx-Rx con-




Fig. 5 Communication and sensing performance versus the number of RF chains across four antenna selection configurations.

figuration demonstrates the most favorable communication-sensing balance. In the communication-centric region with small ρ , it consistently maintains the highest spectral efficiency while gradually improving sensing accuracy as the sensing weight increases. Its “knee point”—the optimal balance point, appears latest among the four configurations, indicating the widest effective trade-off range and the smoothest transition from communication-oriented to sensing-oriented operation. In contrast, the Tx-only configuration exhibits the highest sensitivity to ρ , with its knee point appearing earlier and its spectral efficiency decreasing rapidly once sensing becomes the dominant objective. This behavior indicates a relatively narrow effective trade-off region, where communication performance is sacrificed much earlier in exchange for sensing improvement. The Full-array configuration maintains the best sensing performance over the entire ρ range because of its complete sensing aperture, but its communication performance remains consistently lower than the antenna-selection configurations. The Rx-only configuration shows a smoother transition than the Tx-only configuration; however, its knee point appears earlier and the overall communication-sensing balance remains slightly inferior to that of the Full-array configuration. Overall, these observations further validate the ablation study: transmitter-side antenna selection provides the dominant communication gain, receiver-side antenna selection offers complementary spatial flexibility, and jointly optimizing both sides achieves the best controllable communication-sensing trade-off.

Fig. 7 plots the spectral efficiency against the CRB, which confirms the insights from Fig. 6 and reveals the fundamental communication-sensing trade-off of the four configurations. The curve of the Joint Tx-Rx configuration approaches the Pareto frontier by occupying the upper-left region of the per-




Fig. 6 Communication and sensing performance versus the coefficient ρ across four antenna selection configurations.

formance space. This positioning signifies that for any given sensing accuracy, it delivers the highest possible communication rate, and conversely, for any communication rate, it achieves the best sensing precision. The Tx-only configuration forms the second-best trade-off curve, demonstrating that transmitter-side antenna selection provides the dominant communication gain but still suffers from a noticeable sensing gain because the receiver remains fixed. The Full-array and Rx-only configurations occupy the lower part of the Pareto space and exhibit relatively similar trade-off characteristics. The Full-array configuration achieves a slightly better overall trade-off than the Rx-only configuration because its complete sensing aperture preserves superior sensing accuracy, whereas the communication improvement introduced by receiver-side antenna selection alone is limited. This hierarchy is fundamentally rooted in the management of spatial degrees of freedom. The Full-array configuration provides no spatial adaptation, while the Tx-only and Rx-only configurations introduce flexibility by optimizing only one side of the transceiver. The superiority of the Joint Tx-Rx configuration stems from its end-to-end global resource coordination. By jointly optimizing antenna subsets at both sides, the proposed framework enlarges the achievable communication-sensing performance region and achieves the most efficient Pareto trade-off among the four configurations.

6.3 Scalability and Robustness Evaluation

Fig. 8 illustrates the spectral efficiency and CRB versus the number of transmit antennas while the number of RF chains remains fixed. As N_{TX} increases, the proposed Joint Tx-Rx configuration achieves the most significant SE growth, widening its performance gap over Tx-only and Full-array. This




Fig. 7 Communication-sensing trade-off characteristics across four antenna selection configurations.




Fig. 8 Communication and sensing performance versus the number of transmit antennas across four antenna selection configurations.

trend arises because a larger array provides richer spatial degrees of freedom that the joint transmit and receive selection can effectively exploit, whereas the Full-array configuration quickly saturates due to the fixed RF-chain bottleneck. In terms of sensing, all configurations benefit from the enlarged aperture, but the proposed configuration maintains a CRB only slightly higher than the Full-array configuration while attaining a much higher SE, demonstrating a superior communication-sensing balance across array scales.

Fig. 9 evaluates the spectral efficiency and CRB versus the receive SNR under low-SNR conditions. As observed, the proposed Joint Tx-Rx configuration maintains the highest SE across the entire SNR range, with its advantage over Tx-only becoming more pronounced as the SNR increases. This is because the beamforming and interference suppression gains of joint transmit and receive selection are more effectively




Fig. 9 Communication and sensing performance versus receive SNR across four antenna selection configurations.

realized when the system transitions from a noise-limited to an interference-limited regime. In contrast, the Rx-only and Full-array configurations exhibit noticeably lower SE, confirming that transmit-side selection remains the dominant factor for communication performance even at low SNR. Regarding sensing, the Full-array configuration attains the lowest CRB due to its maximum receive aperture, while the proposed Joint Tx-Rx configuration achieves a CRB only slightly higher, followed by Rx-only and Tx-only. Notably, the CRB of the proposed configuration remains close to that of the Full-array configuration across the entire SNR range, demonstrating its robustness in maintaining high sensing accuracy under low-SNR conditions. These results verify that the proposed joint optimization framework provides a favorable communication-sensing trade-off even in harsh low-SNR environments.

6.4 Comparison with Benchmark Algorithms

To further benchmark the proposed method against classical optimization-based approaches, Fig. 10 compares the proposed Joint Tx-Rx scheme with greedy antenna selection [40], random antenna selection, and the full-array baseline. As observed, the proposed scheme achieves the highest spectral efficiency across the entire power range, followed by the greedy scheme, the full-array scheme, and the random scheme. The greedy scheme surpasses the full-array scheme by concentrating transmit power on a favorable antenna subset, but it still underperforms the proposed method because its antenna selection is independent of the sensing objective. Regarding sensing performance, the full-array scheme exhibits the lowest CRB owing to its maximum receive aperture, but at the cost of severely degraded spectral efficiency. The




Fig. 10 Comparison with greedy, random antenna selection and full-array baselines.

proposed scheme achieves a CRB only slightly higher than the full-array scheme while maintaining a much higher SE, whereas the greedy and random schemes yield significantly higher CRB because their receive selections are not designed for sensing. These results confirm that heuristic or random selection strategies cannot achieve the desired communication–sensing balance, and that joint optimization of transmit and receive antenna subsets is essential.

6.5 Closed-loop Tracking Evaluation

To further evaluate the robustness of the proposed framework under dynamic target movement, Fig. 11 compares the average DOA estimation error and the aggregate spectral efficiency of the closed-loop and open-loop configurations. In this experiment, the time axis is divided into time slots, where each slot corresponds to one complete closed-loop cycle consisting of transmission, echo reception, angle estimation, and feedback-based reconfiguration. The slot duration is assumed to be on the order of the channel coherence time, consistent with Remark 1. The three UAVs initially remain stationary, and then their angles are abruptly shifted at $t = 11$, $t = 16$ and $t = 21$, respectively. The initial angle estimates are assumed to contain a moderate estimation error. As observed, both configurations achieve nearly identical performance before the first sudden movement. When a target suddenly moves, the closed-loop configuration exhibits a temporary performance degradation because the transmit configuration at the current slot is based on the angle estimate from the previous slot. However, as the receiver obtains updated DOA estimates and feeds them back to the transmitter, the system reconfigures the antenna selection and beamformers, allowing the DOA error and spectral efficiency to recover within a few slots. In contrast, the open-loop configuration, which




Fig. 11 Closed-loop versus open-loop tracking performance under sudden target movements.

Table 2 Quantization loss after binary projection.

P_t (dBm)	SE (bit/s/Hz)			CRB (dB)		
	Soft	Binary	Gap (%)	Soft	Binary	Gap (%)
20	23.774	23.772	0.009	-16.396	-16.391	0.032
25	25.306	25.303	0.015	-21.896	-21.884	0.043
30	26.339	26.336	0.014	-26.508	-26.495	0.049

does not update the transmit configuration after the initial slot, suffers from progressively accumulated angle mismatch and cannot recover, resulting in an increasing DOA error and a steadily declining spectral efficiency as the three targets move sequentially. These results demonstrate that the closed-loop sensing-aided reconfiguration mechanism is essential for maintaining reliable ISAC performance under sudden target movements.

6.6 Quantization Loss Evaluation

To quantify the performance loss caused by the final binary projection, Fig. 12 illustrates the spectral efficiency gap and the CRB gap between the relaxed and binary antenna selections as a function of the annealing temperature. As the temperature decreases, the relaxed variables gradually approach exact binary values, leading to a monotonic reduction in both gaps. At the final annealing temperature, the gaps become negligible.

As shown in Fig. 12, the projection loss decreases monotonically with the annealing temperature and becomes negligible at the final temperature. Table 2 further confirms that the spectral efficiency loss is less than 0.02% and the sensing performance loss is below 0.05% across the tested power range. These results indicate that the continuous relaxation provides an essentially exact surrogate for the original mixed-integer




Fig. 12 Quantization loss versus annealing temperature.

antenna selection problem after temperature annealing.

7 Conclusion

In this paper, we proposed a dynamic antenna selection framework for hybrid beamforming in ISAC systems, addressing three fundamental limitations of conventional fixed-subarray architectures through joint transceiver optimization. First, the Gumbel-Softmax relaxation enables adaptive reconfiguration of the antenna subset in response to time-varying channel conditions, overcoming the inability of fixed sub-arrays to adapt to task dynamics. Second, the differentiable optimization under strict cardinality constraints provides the hardware elasticity needed for fault-tolerant and energy-aware operation. Third, the closed-loop sensing-aided beamforming design improves beam quality by aligning the transmit and receive configurations with the estimated UAV angles. Extensive simulations, including ablation studies, scalability and robustness evaluations, and comparisons with benchmark algorithms, demonstrate that the proposed joint design significantly outperforms transmitter-only and full-array schemes, establishing a new and more efficient Pareto frontier for the communication-sensing trade-off. The negligible quantization loss further confirms the effectiveness of the Gumbel-based relaxation for practical deployment. This work confirms that a holistic and adaptive co-design of both the transmitter and receiver is crucial for unlocking the full potential of ISAC in future 6G networks.

Appendix A: CRB Derivation

Define $\boldsymbol{\theta} = [\theta_1, \dots, \theta_K]^T$, $\boldsymbol{\beta}_R = [\beta_{R,1}, \dots, \beta_{R,K}]^T$, $\boldsymbol{\beta}_I = [\beta_{I,1}, \dots, \beta_{I,K}]^T$, where $\beta_{R,k} = \text{Re}(\beta_k)$ and $\beta_{I,k} = \text{Im}(\beta_k)$. The $3K$ real-valued parameters to be

estimated are given by $\boldsymbol{\xi} = [\boldsymbol{\theta}^T, \boldsymbol{\beta}_R^T, \boldsymbol{\beta}_I^T]^T$. The received signal vector $\mathbf{r}$ follows a complex Gaussian distribution, denoted as $\mathbf{r} \sim \mathcal{CN}(\boldsymbol{\mu}(\boldsymbol{\xi}), \mathbf{C})$, where $\boldsymbol{\mu}(\boldsymbol{\xi}) = \sum_k^K \beta_k \mathbf{b}(\theta_k) \mathbf{a}^H(\theta_k) \tilde{\mathbf{F}}_{\text{RF}} \mathbf{F}_{\text{BBS}}$ and $\mathbf{C} = \mathbb{E}[\tilde{\mathbf{n}}\tilde{\mathbf{n}}^H] = \sigma_s^2 \cdot \text{diag}(\boldsymbol{\omega})$. The noise covariance matrix $\mathbf{C}$ is independent of the parameters, so the (i, j) -th element of the Fisher information matrix (FIM) is given by

$$\mathbf{J}_{i,j} = 2\text{Re} \left\{ \frac{\partial \boldsymbol{\mu}^H(\boldsymbol{\xi})}{\partial \xi_i} \mathbf{C}^{-1} \frac{\partial \boldsymbol{\mu}(\boldsymbol{\xi})}{\partial \xi_j} \right\}. \quad (\text{A1})$$

Then, the FIM is partitioned into a 3×3 block matrix with each block being a $K \times K$ matrix:

$$\mathbf{J} = \begin{bmatrix} \mathbf{J}_{\boldsymbol{\theta}\boldsymbol{\theta}} & \mathbf{J}_{\boldsymbol{\theta}\boldsymbol{\beta}_R} & \mathbf{J}_{\boldsymbol{\theta}\boldsymbol{\beta}_I} \\ \mathbf{J}_{\boldsymbol{\beta}_R\boldsymbol{\theta}} & \mathbf{J}_{\boldsymbol{\beta}_R\boldsymbol{\beta}_R} & \mathbf{J}_{\boldsymbol{\beta}_R\boldsymbol{\beta}_I} \\ \mathbf{J}_{\boldsymbol{\beta}_I\boldsymbol{\theta}} & \mathbf{J}_{\boldsymbol{\beta}_I\boldsymbol{\beta}_R} & \mathbf{J}_{\boldsymbol{\beta}_I\boldsymbol{\beta}_I} \end{bmatrix}. \quad (\text{A2})$$

The CRB matrix for $\boldsymbol{\theta}$ is given by the Schur complement as [41]

$$\text{CRB}(\boldsymbol{\theta}) = \left(\mathbf{J}_{\boldsymbol{\theta}\boldsymbol{\theta}} - \mathbf{J}_{\boldsymbol{\theta}\boldsymbol{\beta}} \mathbf{J}_{\boldsymbol{\beta}\boldsymbol{\beta}}^{-1} \mathbf{J}_{\boldsymbol{\beta}\boldsymbol{\theta}} \right)^{-1}, \quad (\text{A3})$$

where $\mathbf{J}_{\boldsymbol{\theta}\boldsymbol{\beta}} = [\mathbf{J}_{\boldsymbol{\theta}\boldsymbol{\beta}_R}, \mathbf{J}_{\boldsymbol{\theta}\boldsymbol{\beta}_I}]$, $\mathbf{J}_{\boldsymbol{\beta}\boldsymbol{\theta}} = \mathbf{J}_{\boldsymbol{\theta}\boldsymbol{\beta}}^T$ and

$$\mathbf{J}_{\boldsymbol{\beta}\boldsymbol{\beta}} = \begin{bmatrix} \mathbf{J}_{\boldsymbol{\beta}_R\boldsymbol{\beta}_R} & \mathbf{J}_{\boldsymbol{\beta}_R\boldsymbol{\beta}_I} \\ \mathbf{J}_{\boldsymbol{\beta}_I\boldsymbol{\beta}_R} & \mathbf{J}_{\boldsymbol{\beta}_I\boldsymbol{\beta}_I} \end{bmatrix}. \quad (\text{A4})$$

Eq. (A3) exhibits a complex structure of the CRB matrix. To simplify the resulting expression, we assume that the targets are sufficiently separated from each other. Owing to this spatial separation, the inter-target coupling tends to zero. Consequently, the FIM can be approximated by a block-diagonal form as

$$\mathbf{J} \approx \text{blkdiag}(\mathbf{J}_{\boldsymbol{\eta}_1 \boldsymbol{\eta}_1}, \dots, \mathbf{J}_{\boldsymbol{\eta}_K \boldsymbol{\eta}_K}), \quad (\text{A5})$$

where $\boldsymbol{\eta}_k = [\theta_k, \beta_{R,k}, \beta_{I,k}]^T$. The inverse of the FIM is also block-diagonal, expressed as

$$\mathbf{J}^{-1} = \text{blkdiag}(\mathbf{J}_{\boldsymbol{\eta}_1 \boldsymbol{\eta}_1}^{-1}, \dots, \mathbf{J}_{\boldsymbol{\eta}_K \boldsymbol{\eta}_K}^{-1}). \quad (\text{A6})$$

Therefore, the CRB of the k -th UAV angle θ_k is given by

$$\text{CRB}(\theta_k) = [\mathbf{J}_{\boldsymbol{\eta}_k \boldsymbol{\eta}_k}^{-1}]_{1,1} = \left(\mathbf{J}_{\boldsymbol{\theta}_k \boldsymbol{\theta}_k} - \mathbf{J}_{\boldsymbol{\theta}_k \boldsymbol{\beta}_k} \mathbf{J}_{\boldsymbol{\beta}_k \boldsymbol{\beta}_k}^{-1} \mathbf{J}_{\boldsymbol{\beta}_k \boldsymbol{\theta}_k} \right)^{-1}, \quad (\text{A7})$$

where $\mathbf{J}_{\boldsymbol{\theta}_k \boldsymbol{\beta}_k} = [\mathbf{J}_{\boldsymbol{\theta}_k \boldsymbol{\beta}_{R,k}}, \mathbf{J}_{\boldsymbol{\theta}_k \boldsymbol{\beta}_{I,k}}]$, $\mathbf{J}_{\boldsymbol{\beta}_k \boldsymbol{\theta}_k} = \mathbf{J}_{\boldsymbol{\theta}_k \boldsymbol{\beta}_k}^T$ and

$$\mathbf{J}_{\boldsymbol{\beta}_k \boldsymbol{\beta}_k} = \begin{bmatrix} \mathbf{J}_{\boldsymbol{\beta}_{R,k} \boldsymbol{\beta}_{R,k}} & \mathbf{J}_{\boldsymbol{\beta}_{R,k} \boldsymbol{\beta}_{I,k}} \\ \mathbf{J}_{\boldsymbol{\beta}_{I,k} \boldsymbol{\beta}_{R,k}} & \mathbf{J}_{\boldsymbol{\beta}_{I,k} \boldsymbol{\beta}_{I,k}} \end{bmatrix}. \quad (\text{A8})$$

Under high-SNR conditions, β_k can be estimated with high accuracy. This results in the exceedingly large Fisher information $\mathbf{J}_{\boldsymbol{\beta}_k \boldsymbol{\beta}_k}$, thereby rendering the influence of the performance penalty term $\mathbf{J}_{\boldsymbol{\theta}_k \boldsymbol{\beta}_k} \mathbf{J}_{\boldsymbol{\beta}_k \boldsymbol{\beta}_k}^{-1} \mathbf{J}_{\boldsymbol{\beta}_k \boldsymbol{\theta}_k}$ negligible. Therefore,

a tractable expression of the CRB is derived as

$$\begin{aligned} \text{CRB}(\theta_k) = & \frac{\sigma_s^2}{2|\beta_k|^2} \left[\dot{\mathbf{b}}^H(\theta_k) \dot{\mathbf{b}}(\theta_k) \sum_{l=1}^K \left| \mathbf{a}^H(\theta_k) \tilde{\mathbf{F}}_{\text{RF}} \mathbf{v}_l \right|^2 \right. \\ & + \mathbf{b}^H(\theta_k) \mathbf{b}(\theta_k) \sum_{l=1}^K \left| \dot{\mathbf{a}}^H(\theta_k) \tilde{\mathbf{F}}_{\text{RF}} \mathbf{v}_l \right|^2 \\ & \left. + 2\text{Re} \left\{ \dot{\mathbf{b}}^H(\theta_k) \mathbf{b}(\theta_k) \sum_{l=1}^K \dot{\mathbf{a}}^H(\theta_k) \tilde{\mathbf{F}}_{\text{RF}} \mathbf{v}_l \mathbf{v}_l^H \tilde{\mathbf{F}}_{\text{RF}}^H \mathbf{a}(\theta_k) \right\} \right]^{-1}. \end{aligned} \quad (\text{A9})$$

It is satisfied that $\dot{\mathbf{b}}^H(\theta) \mathbf{b}(\theta) = 0$ when the array origin is at the array center, so the CRB expression (A9) is simplified as

$$\begin{aligned} \text{CRB}(\theta_k) = & \frac{\sigma_s^2}{2|\beta_k|^2} \left[\dot{\mathbf{b}}^H(\theta_k) \dot{\mathbf{b}}(\theta_k) \sum_{l=1}^K \left| \mathbf{a}^H(\theta_k) \tilde{\mathbf{F}}_{\text{RF}} \mathbf{v}_l \right|^2 \right. \\ & \left. + \mathbf{b}^H(\theta_k) \mathbf{b}(\theta_k) \sum_{l=1}^K \left| \dot{\mathbf{a}}^H(\theta_k) \tilde{\mathbf{F}}_{\text{RF}} \mathbf{v}_l \right|^2 \right]^{-1}. \end{aligned} \quad (\text{A10})$$

Appendix B: Proof of Theorem 1

We proceed by analyzing the total change in the objective function between two consecutive iterations, from t to $t+1$. We decompose this change by introducing an intermediate term:

$$\begin{aligned} \Delta \mathcal{L}^{(t)} &= \mathcal{L}(\boldsymbol{\Theta}^{(t+1)}; \tau^{(t+1)}) - \mathcal{L}(\boldsymbol{\Theta}^{(t)}; \tau^{(t)}) \\ &= \left[\mathcal{L}(\boldsymbol{\Theta}^{(t+1)}; \tau^{(t+1)}) - \mathcal{L}(\boldsymbol{\Theta}^{(t+1)}; \tau^{(t)}) \right] \\ &\quad + \left[\mathcal{L}(\boldsymbol{\Theta}^{(t+1)}; \tau^{(t)}) - \mathcal{L}(\boldsymbol{\Theta}^{(t)}; \tau^{(t)}) \right]. \end{aligned} \quad (\text{B1})$$

Let us analyze the two bracketed terms. Define $\Delta \mathcal{L}_{\text{opt}}^{(t)} = \mathcal{L}(\boldsymbol{\Theta}^{(t+1)}; \tau^{(t)}) - \mathcal{L}(\boldsymbol{\Theta}^{(t)}; \tau^{(t)})$ as the ascent term from the optimization step at a fixed temperature $\tau^{(t)}$. Define $\Delta \mathcal{L}_{\tau}^{(t)} = \mathcal{L}(\boldsymbol{\Theta}^{(t+1)}; \tau^{(t+1)}) - \mathcal{L}(\boldsymbol{\Theta}^{(t+1)}; \tau^{(t)})$ as the perturbation term due to the change in temperature.

First, we analyze the ascent term $\Delta \mathcal{L}_{\text{opt}}^{(t)}$. By the definition of the AO algorithm, $\boldsymbol{\Theta}^{(t+1)}$ is obtained by applying block coordinate ascent to $\boldsymbol{\Theta}^{(t)}$. Each block-wise update is guaranteed to not decrease the objective function value. Therefore, after a full cycle of updates, we have:

$$\Delta \mathcal{L}_{\text{opt}}^{(t)} = \mathcal{L}(\boldsymbol{\Theta}^{(t+1)}; \tau^{(t)}) - \mathcal{L}(\boldsymbol{\Theta}^{(t)}; \tau^{(t)}) \geq 0. \quad (\text{B2})$$

Next, we bound the perturbation term $\Delta \mathcal{L}_{\tau}^{(t)}$. Using Property 2 (Lipschitz continuity), we can bound its absolute value as:

$$\begin{aligned} \left| \Delta \mathcal{L}_{\tau}^{(t)} \right| &= \left| \mathcal{L}(\boldsymbol{\Theta}^{(t+1)}; \tau^{(t+1)}) - \mathcal{L}(\boldsymbol{\Theta}^{(t+1)}; \tau^{(t)}) \right| \\ &\leq K_{\tau} \left| \tau^{(t+1)} - \tau^{(t)} \right|. \end{aligned} \quad (\text{B3})$$

Now, let us consider the sum of the ascent terms from $t = 0$ to $N - 1$. By rearranging (B1), we have $\Delta \mathcal{L}_{\text{opt}}^{(t)} = \mathcal{L}(\boldsymbol{\Theta}^{(t+1)}; \tau^{(t+1)}) - \mathcal{L}(\boldsymbol{\Theta}^{(t)}; \tau^{(t)}) - \Delta \mathcal{L}_{\tau}^{(t)}$. Summing this over t :

$$\begin{aligned} \sum_{t=0}^{N-1} \Delta \mathcal{L}_{\text{opt}}^{(t)} &= \sum_{t=0}^{N-1} \left[\mathcal{L}(\boldsymbol{\Theta}^{(t+1)}; \tau^{(t+1)}) - \mathcal{L}(\boldsymbol{\Theta}^{(t)}; \tau^{(t)}) \right] \\ &\quad - \sum_{t=0}^{N-1} \Delta \mathcal{L}_{\tau}^{(t)} = \mathcal{L}(\boldsymbol{\Theta}^{(N)}; \tau^{(N)}) - \mathcal{L}(\boldsymbol{\Theta}^{(0)}; \tau^{(0)}) - \sum_{t=0}^{N-1} \Delta \mathcal{L}_{\tau}^{(t)}. \end{aligned} \quad (\text{B4})$$

Since $\Delta \mathcal{L}_{\text{opt}}^{(t)} \geq 0$, its sum is also non-negative. We can bound this sum:

$$\begin{aligned} \sum_{t=0}^{N-1} \Delta \mathcal{L}_{\text{opt}}^{(t)} &\leq \left| \mathcal{L}(\boldsymbol{\Theta}^{(N)}; \tau^{(N)}) \right| \\ &\quad + \left| \mathcal{L}(\boldsymbol{\Theta}^{(0)}; \tau^{(0)}) \right| + \sum_{t=0}^{N-1} \left| \Delta \mathcal{L}_{\tau}^{(t)} \right|. \end{aligned} \quad (\text{B5})$$

By Property 1, $\left| \mathcal{L}(\boldsymbol{\Theta}^{(N)}; \tau^{(N)}) \right|$ and $\left| \mathcal{L}(\boldsymbol{\Theta}^{(0)}; \tau^{(0)}) \right|$ are bounded by L_{sup} . Using the bound from (B3) for the final term, we get:

$$\sum_{t=0}^{N-1} \Delta \mathcal{L}_{\text{opt}}^{(t)} \leq 2L_{\text{sup}} + K_{\tau} \sum_{t=0}^{N-1} \left| \tau^{(t+1)} - \tau^{(t)} \right|. \quad (\text{B6})$$

As $N \rightarrow \infty$, the right-hand side remains finite due to Property 3, which states that $\sum_{t=0}^{N-1} \left| \tau^{(t+1)} - \tau^{(t)} \right|$ converges. Since $\{\Delta \mathcal{L}_{\text{opt}}^{(t)}\}$ is a sequence of non-negative terms and its infinite sum is bounded, the series $\sum_{t=0}^{N-1} \Delta \mathcal{L}_{\text{opt}}^{(t)}$ must converge. A necessary condition for the convergence of a series of non-negative terms is that the terms themselves must converge to zero. Therefore,

$$\begin{aligned} \lim_{t \rightarrow \infty} \Delta \mathcal{L}_{\text{opt}}^{(t)} &= \lim_{t \rightarrow \infty} \left[\mathcal{L}(A(\boldsymbol{\Theta}^{(t)}, \tau^{(t)}); \tau^{(t)}) - \mathcal{L}(\boldsymbol{\Theta}^{(t)}; \tau^{(t)}) \right] = 0, \end{aligned} \quad (\text{B7})$$

where $A(\boldsymbol{\Theta}^{(t)}, \tau^{(t)})$ represents the operation of performing one cycle of block coordinate ascent on $\boldsymbol{\Theta}^{(t)}$ at $\tau^{(t)}$.

This result implies that as $t \rightarrow \infty$, the improvement gained from one full iteration of the AO algorithm vanishes. Let $\boldsymbol{\Theta}^*$ be any limit point of the sequence $\{\boldsymbol{\Theta}^{(t)}\}$. Since when $t \rightarrow \infty$, $\tau^{(t)} \rightarrow 0$, any such limit point must satisfy:

$$\mathcal{L}(A(\boldsymbol{\Theta}^*, 0); 0) - \mathcal{L}(\boldsymbol{\Theta}^*; 0) = 0. \quad (\text{B8})$$

Eq. (B8) states that at $\boldsymbol{\Theta}^*$, no further improvement can be made by applying the block coordinate ascent procedure. This concludes the proof that the algorithm converges to a stationary point of the desired problem.

Declaration of competing interest

The authors have no competing interests to declare that are relevant to the content of this article.

Reference

- [1] F. Liu, Y. Cui, C. Masouros, J. Xu, T. X. Han, Y. C. Eldar, and S. Buzzi, Integrated sensing and communications: Toward dual-functional wireless networks for 6G and beyond, *IEEE J. Sel. Areas Commun.*, vol. 40, no. 6, pp. 1728–1767, Jun. 2022.
- [2] A. Cengiz and Y. Kabalci, Towards unified wireless systems: ISAC technologies in the 6G era, *2025 7th Global Power, Energy and Communication Conference (GPECOM)*, Bochum, Germany, 2025, pp. 1026–1031.
- [3] D. K. Pin Tan *et al.*, Integrated sensing and communication in 6G: Motivations, use cases, requirements, challenges and future directions, *2021 1st IEEE International Online Symposium on Joint Communications & Sensing (JC&S)*, Dresden, Germany, 2021, pp. 1–6.
- [4] N. González-Prelcic *et al.*, The integrated sensing and communication revolution for 6G: Vision, techniques, and applications, *Proc. IEEE*, vol. 112, no. 7, pp. 676–723, Jul. 2024.
- [5] B. Li, Y. Zou, T. Wu, Z. Zhang, M. Chen and Y. Jiang, Security and reliability tradeoff of NOMA based hybrid satellite-terrestrial network with a friendly jammer, *IEEE Trans. Veh. Technol.*, vol. 74, no. 2, pp. 3439–3444, Feb. 2025.
- [6] Y. Wang *et al.*, ISAC enabled cooperative detection for cellular-connected UAV network, *IEEE Trans. Wireless Commun.*, vol. 24, no. 2, pp. 1541–1554, Feb. 2025.
- [7] Y. Li *et al.*, RIS-based physical layer security for integrated sensing and communication: A comprehensive survey, *IEEE Internet Things J.*, vol. 12, no. 16, pp. 32444–32468, Aug. 15, 2025.
- [8] J. A. Zhang *et al.*, An overview of signal processing techniques for joint communication and radar sensing, *IEEE J. Sel. Top. Signal Process.*, vol. 15, no. 6, pp. 1295–1315, Nov. 2021.
- [9] Y. Cui, F. Liu, X. Jing, and J. Mu, Integrating sensing and communications for ubiquitous IoT: Applications, trends, and challenges, *IEEE Netw.*, vol. 36, no. 5, pp. 158–167, Sep./Oct. 2022.
- [10] F. Liu, L. Zhou, C. Masouros, A. Li, W. Luo, and A. Petropulu, Toward dual-functional radar-communication systems: Optimal waveform design, *IEEE Trans. Signal Process.*, vol. 66, no. 16, pp. 4264–4279, Aug. 2018.
- [11] R. W. Heath, N. Gonzalez-Prelcic, S. Rangan, W. Roh, and A. M. Sayeed, An overview of signal processing techniques for millimeter wave MIMO systems, *IEEE J. Sel. Top. Signal Process.*, vol. 10, no. 3, pp. 436–453, Apr. 2016.
- [12] J. A. Zhang *et al.*, Enabling joint communication and radar sensing in mobile networks—A survey, *IEEE Commun. Surv. Tuts.*, vol. 24, no. 1, pp. 306–345, Firstquarter 2022.
- [13] A. F. Molisch *et al.*, Hybrid beamforming for massive MIMO: A Survey, *IEEE Commun. Mag.*, vol. 55, no. 9, pp. 134–141, Sep. 2017.
- [14] F. Liu, C. Masouros, A. P. Petropulu, H. Griffiths, and L. Hanzo, Joint radar and communication design: Applications, state-of-the-art, and the road ahead, *IEEE Trans. Commun.*, vol. 68, no. 6, pp. 3834–3862, Jun. 2020.
- [15] C. Qi, W. Ci, J. Zhang and X. You, Hybrid beamforming for millimeter wave MIMO integrated sensing and communications, *IEEE Commun. Lett.*, vol. 26, no. 5, pp. 1136–1140, May 2022.
- [16] D. N. Dao, H. Zhang, A. B. J. Kokkeler and Y. Miao, Joint beamforming for multi-user multi-target FD ISAC system: A hybrid GRQ-GA approach, *2025 IEEE Wireless Communications and Networking Conference (WCNC)*, Milan, Italy, 2025, pp. 1–6.
- [17] S. Park, A. Alkhateeb and R. W. Heath, Dynamic subarrays for hybrid precoding in wideband mmwave MIMO systems, *IEEE Trans. Wireless Commun.*, vol. 16, no. 5, pp. 2907–2920, May 2017.
- [18] J. Jiang, Y. Yuan and L. Zhen, Multi-user hybrid precoding for dynamic subarrays in mmwave massive MIMO systems, *IEEE Access*, vol. 7, pp. 101718–101728, 2019.
- [19] W. Nie, M. Liu, J. Chen, W. Tan and C. Li, Spectrum and energy efficiency of massive MIMO for hybrid architectures with phase shifter and switches in IoT networks, *IEEE Internet Things J.*, vol. 11, no. 6, pp. 9656–9667, 15 Mar., 2024.
- [20] C. Qi, Q. Liu, X. Yu and G. Y. Li, Hybrid precoding for mixture use of phase shifters and switches in mmWave massive MIMO, *IEEE Trans. Commun.*, vol. 70, no. 6, pp. 4121–4133, Jun. 2022.
- [21] H. Nosrati, E. Aboutanios, D. Smith and X. Wang, Switch-based hybrid precoding in mmWave massive MIMO systems, *2019 27th European Signal Processing Conference (EUSIPCO)*, A Coruna, Spain, 2019, pp. 1–5.
- [22] S. Asaad, A. Bereyhi, A. M. Rabiei, R. R. Müller and R. F. Schaefer, Optimal transmit antenna selection for massive MIMO wiretap channels, *IEEE J. Sel. Areas Commun.*, vol. 36, no. 4, pp. 817–828, Apr. 2018.
- [23] F. Shu, X. Jiang, X. Liu, L. Xu, G. Xia and J. Wang, Precoding and transmit antenna subarray selection for secure hybrid spatial modulation, *IEEE Trans. Wireless Commun.*, vol. 20, no. 3, pp. 1903–1917, Mar. 2021.
- [24] J. Yu, S. Zhao, M. Tian, Q. Li and D. Tang, Joint transmit antenna selection and beamforming design for millimeter wave massive MIMO NOMA secure networks, *IEEE Commun. Lett.*, vol. 28, no. 10, pp. 2407–2411, Oct. 2024.
- [25] T. Ding, Y. Zhao, L. Li, D. Hu and L. Zhang, Hybrid precoding for beamspace MIMO systems with sub-connected switches: A machine learning approach, *IEEE Access*, vol. 7, pp. 143273–143281, 2019.
- [26] Y. Hei, C. Liu, W. Li, L. Ma and M. Lan, CNN based hybrid precoding for mmwave MIMO systems with adaptive switching



- module and phase modulation array, *IEEE Trans. Wireless Commun.*, vol. 21, no. 12, pp. 10489-10501, Dec. 2022.
- [27] W. Khawaja, I. Guvenc, D. W. Matolak, U. -C. Fiebig and N. Schneckenburger, A survey of air-to-ground propagation channel modeling for unmanned aerial vehicles, *IEEE Commun. Surv. Tutorials*, vol. 21, no. 3, pp. 2361-2391, Thirdquarter 2019.
- [28] A. A. Khuwaja, Y. Chen, N. Zhao, M. -S. Alouini and P. Dobbins, A survey of channel modeling for UAV communications, *IEEE Commun. Surv. Tutorials*, vol. 20, no. 4, pp. 2804-2821, Fourthquarter 2018.
- [29] F. Liu and C. Masouros, Hybrid beamforming with sub-arrayed MIMO radar: enabling joint sensing and communication at mmWave band, *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, May 2019, pp. 7770-7774.
- [30] F. Liu, W. Yuan, C. Masouros and J. Yuan, Radar-assisted predictive beamforming for vehicular links: Communication served by sensing, *IEEE Trans. Wireless Commun.*, vol. 19, no. 11, pp. 7704-7719, Nov. 2020.
- [31] M. Steer, *Microwave and RF Design IV: Modules*. Raleigh, NC, USA: SciTech Publishing, 2016.
- [32] P. Xie, Y. Wan, T. Li and R. Ju, Low complexity channel estimation algorithm using paired spatial signatures for UAV 3D MIMO systems, *J Wireless Com Network*, vol. 2021, no. 53, Mar. 2021.
- [33] E. Jang, S. Gu, and B. Poole, Categorical reparameterization with Gumbel-Softmax, *arXiv preprint arXiv:1611.01144*, 2017.
- [34] C. J. Maddison, A. Mnih, and Y. W. Teh, The concrete distribution: A continuous relaxation of discrete random variables, *arXiv preprint arXiv:1611.00712*, 2017.
- [35] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. Cambridge, MA, USA: MIT Press, 2016.
- [36] L. Franceschi, M. Donini, P. Frasconi, and M. Pontil, Forward and reverse gradient-based hyperparameter optimization, in *Proc. 34th Int. Conf. Mach. Learn. (ICML)*, Sydney, Australia, 2017, pp. 1165-1173.
- [37] K. Shen, and W. Yu, Fractional programming for communication systems—part I: Power control and beamforming, *IEEE Trans. Signal Process.*, vol. 66, no. 10, pp. 2616-2630, May 2018.
- [38] K. Shen, and W. Yu, Fractional programming for communication systems—part II: Uplink scheduling via matching, *IEEE Trans. Signal Process.*, vol. 66, no. 10, pp. 2631-2644, May 2018.
- [39] R. Schmidt, Multiple emitter location and signal parameter estimation, *IEEE Trans. Antennas Propag.*, vol. 34, no. 3, pp. 276-280, Mar. 1986.
- [40] X. Li, Y. Lin, X. Yang, W. Pu, M. Meng and Y. Hei, Greedy hybrid beamforming for multiuser MmWave MIMO systems, *IET Commun.*, vol. 1, pp. 2800-2805, Dec. 2017.
- [41] S. M. Kay, *Fundamentals of Statistical Signal Processing:*

Estimation Theory. Englewood Cliffs, NJ, USA: Prentice Hall, 1993.

Author biography



Chenghao Zhu received the B.S. degree from Southeast University, Nanjing, China, in 2024. He is currently pursuing the Ph.D. degree with the School of Information Science and Engineering, Southeast University, Nanjing. His current research interests include hybrid beamforming, integrated sensing and communication, dynamic antenna systems, and near-field communications.



Liang Wu received the B.S., M.S., and Ph.D. degrees from the School of Information Science and Engineering, Southeast University, Nanjing, China, in 2007, 2010, and 2013, respectively. From September 2011 to March 2013, he was with the School of Electrical Engineering and Computer Science, Oregon State University, as a Visiting Ph.D. Student. In September 2013, he joined the National Mobile Communications Research Laboratory, Frontiers Science Center for Mobile Information Communication and Security, Southeast University. Since April 2024, he has been a full professor in Southeast University. He is also with Purple Mountain Laboratories, Nanjing. His research interests include multiple-antenna techniques, interference coordination, optical wireless communication systems, and wireless localization systems.