

Nuisance-aware geometry-reconfigurable cell-free ISAC: identifiability, joint design, and calibration limits



Changhao He

Computer, Electrical and Mathematical Sciences and Engineering (CEMSE) Division, King Abdullah University of Science and Technology (KAUST), Thuwal, Saudi Arabia

E-mail: changhao.he@kaust.edu.sa.

Highlights:

- Nuisance-aware analysis reveals zero transmit-angle information under rank-one probing.
- Joint covariance and array-shape optimization improves localization accuracy.
- Residual calibration errors limit the gains from geometry reconfiguration.

Abstract: Geometry-reconfigurable arrays offer additional spatial flexibility for cell-free integrated sensing and communication (ISAC), but unknown complex bistatic gains can limit their localization benefit. We derive a nuisance-aware equivalent Fisher information matrix that isolates identifiable angular information after gain elimination. The analysis shows that rank-one probing provides no transmit-angle information, even when a known-gain model predicts a positive value; rank at least two is necessary but not sufficient. Mapping angular information to position separates local aperture sensitivity from network angular diversity. We then jointly optimize communication and probing covariances and transmit/receive array shapes using semidefinite covariance optimization and safeguarded shape search, with deterministic rank-one recovery of communication covariances. A hybrid information model incorporates residual orientation and actuator errors to characterize calibration limits. Under a physical 28-GHz model, joint reconfiguration achieves a 5.51-dB paired Cramér-Rao bound gain over collapsed-geometry covariance-only design across 48 trials. Maximum-likelihood localization remains within 0.45 dB of the root-mean-square position error bound at the tested efficiency points. Orientation uncertainty substantially degrades localization, while covariance robustification offers little improvement in the evaluated setting. These results identify when geometry reconfiguration yields localization gains and how residual calibration errors limit their realization.

Keywords: cell-free ISAC; equivalent Fisher information; geometry-reconfigurable array; movable antenna; Cramér-Rao bound; identifiability; array calibration; variable-projection maximum likelihood



Copyright©2026 by the author. Published by ELSP. This work is licensed under a Creative Commons Attribution 4.0 International License, which permits unrestricted use, distribution, and reproduction in any medium provided the original work is properly cited.

1. Introduction

Cell-free integrated sensing and communication (ISAC) uses distributed access points (APs) to serve users and observe targets from multiple directions. Its localization accuracy depends on both local angular information and network viewing geometry. Geometry-reconfigurable arrays add a controllable sensing resource by allowing radiating elements to move within prescribed regions.

However, aperture changes do not necessarily produce identifiable information. When a bistatic return contains an unknown complex gain, steering variations parallel to the received signal are indistinguishable from gain changes. A design based on raw steering derivatives can overestimate the localization benefit of reconfiguration. Residual orientation and actuator errors impose a further limitation: the information predicted for a commanded geometry may not be available when that geometry is imperfectly calibrated.

This paper studies which geometry-induced variations remain identifiable, how they contribute to position estimation, and how to exploit them under shared communication and sensing resources.

The main contributions are as follows:

- **Identifiable sensing information.** We derive a nuisance-aware equivalent Fisher information matrix (EFIM) by eliminating unknown bistatic complex gains. The result characterizes the angular information retained by each aperture and its mapping to position. In particular, rank-one probing provides no transmit-angle information; rank at least two is necessary but not sufficient. The position EFIM separates local information intensity from network angular diversity and gives the condition for a finite Cramér–Rao bound (CRB).
- **Joint covariance and geometry design.** We optimize communication and dedicated-probing covariances together with transmit/receive array shapes under user-centric serving constraints and per-AP power budgets. For fixed shapes, the rank-relaxed covariance problem is a semidefinite program on the finite-CRB domain and admits deterministic rank-one recovery of communication covariances. The safeguarded shape search terminates at a coordinatewise endpoint-stable design over the prescribed finite endpoint set.
- **Calibration limits and numerical validation.** We incorporate orientation and actuator errors into the complete hybrid information model and derive an explicit orientation information ceiling. Numerical experiments compare random geometries, test redesign across target locations, and evaluate finite-sample maximum-likelihood localization. Joint reconfiguration achieves a 5.51-dB paired CRB gain over collapsed-geometry covariance-only design in the physical 28-GHz campaign, while calibration sweeps identify a regime in which covariance reallocation offers little improvement.

2. Related work and positioning

2.1. Cell-free and multistatic ISAC

Cell-free ISAC combines coordinated communication with spatially distributed sensing. Existing studies optimize beamforming, detection, and power allocation, demonstrating the value of multiple bistatic views [1–3]. The information analysis developed here also applies to multistatic sensing networks with the same observation model. Its cell-free ISAC role lies in the joint resource design: user-centric

communication and dedicated probing share per-AP power budgets, and probing contributes interference to communication users.

2.2. Movable and fluid antennas for ISAC

Fluid-antenna research establishes diversity and port-selection principles [4–6], while movable-antenna studies develop position-dependent channel models and array-design methods [7–9]. Sensing and ISAC extensions jointly optimize antenna positions and beamforming [10–13]. Recent work also considers synchronization errors, hardware distortion, and double-side fluid antennas in distributed systems [14–16]. In our earlier CRB analysis of cell-free ISAC with fluid intelligent metasurfaces [17], the bistatic complex gains were treated as known. The present work focuses on the geometry-induced information that survives unknown-gain elimination and its implications for joint design and calibration.

2.3. Nuisance-aware bounds and array calibration

Schur-complement elimination of nuisance parameters and hybrid information bounds are established estimation tools [18–20]. Distributed radar and network-localization studies characterize the role of viewing geometry [21–23], while array self-calibration studies analyze uncertain sensor locations and their identifiability limits [24–26]. Here, the commanded aperture is itself a design variable, and residual calibration uncertainty limits the gains attainable through reconfiguration. We also evaluate maximum-likelihood localization to distinguish information-bound improvements from finite-sample estimator performance, a distinction emphasized in recent localization research [27].

Table 1 collects the principal notation used throughout the formulation and numerical evaluation.

Table 1. Compact nomenclature used throughout the revised manuscript. Bold lower/upper-case symbols denote vectors/matrices; nonbold symbols are scalars.

Symbol	Meaning
t, r, k, n, s	Tx AP, Rx AP, communication-user, array-element, and calibration-scenario indices
$M_t, M_r, K_c, N_m, L, N_{sc}$	numbers of Tx APs, Rx APs, users, elements on array m , sensing snapshots, and calibration scenarios
$\mathbf{p}, \mathbf{z}_m, \tilde{\mathbf{z}}_m$	target position, commanded element-displacement vector, and displacement envelope
$\mathbf{a}_m, \dot{\mathbf{a}}_m, \mathbf{D}_m$	steering vector, angular derivative, and diagonal phase-rate matrix
$\mathbf{Q}_k, \mathbf{S}_t, \mathbf{S}, \mathbf{D}_k^{\text{srv}}$	user- k communication covariance, Tx- t probing covariance, stacked probing covariance, and fixed serving mask
$\Gamma_k, P_t, q_t, d_{\min, m}$	SINR target, Tx-AP power budget, illumination floor, and minimum element spacing
$q_t, \kappa_t, \mathbf{I}_m, \mathbf{w}_m$	scalar illumination, nuisance-aware Tx angular term, local bearing information, and position-EFIM weight
$\boldsymbol{\gamma}, \boldsymbol{\delta}_m, \mathbf{e}_m, \mathbf{P}_m^\perp$	calibration vector, orientation error, actuator error, and common-displacement-removing projector
$C_H, c_s, \text{CVaR}_{95}$	hybrid CRB, scenario-wise conditional/local CRB, and 95%-level conditional value-at-risk (CVaR) tail-risk statistic

3. System model

Consider M_t transmitting APs and M_r receiving APs with known positions $\mathbf{p}_t^{\text{Tx}}, \mathbf{p}_r^{\text{Rx}} \in \mathbb{R}^2$. A coarse acquisition or previous tracking stage supplies a nominal target location $\mathbf{p} = [x, y]^T$ at which the local information design is formed. Figure 1 summarizes the physical cell-free architecture and the joint geometry/covariance design loop used throughout the paper.




Figure 1. Geometry-reconfigurable cell-free ISAC. **(a)** Distributed transmit APs serve users and illuminate a target, while receive APs collect bistatic returns. Open and filled markers indicate reference and displaced element positions; δ and e denote residual orientation and actuator errors, respectively; **(b)** Unknown-gain elimination yields the position EFIM used for alternating covariance and array-shape optimization. Residual calibration errors enter the hybrid information extension, and the resulting design updates are returned to the APs through the central processor.

For a generic N_m -element one-dimensional reconfigurable array, let $\mathbf{z}_m = [z_{m,1}, \dots, z_{m,N_m}]^T$ denote element displacements with $|z_{m,n}| \leq \tilde{z}_m$. The normalized response is

$$[\mathbf{a}_m(\theta_m, \mathbf{z}_m)]_n = \frac{1}{\sqrt{N_m}} e^{j\phi_{m,n}}, \quad (1)$$

where

$$\phi_{m,n} = (n-1)\pi \sin \theta_m + \frac{2\pi}{\lambda} z_{m,n} \cos \theta_m. \quad (2)$$

Define

$$d_{m,n} = (n-1)\pi \cos \theta_m - \frac{2\pi}{\lambda} z_{m,n} \sin \theta_m, \quad (3)$$

$$\dot{\mathbf{a}}_m = j\mathbf{D}_m \mathbf{a}_m, \quad \mathbf{D}_m = \text{diag}(\mathbf{d}_m). \quad (4)$$

The nominal coordinates lie on one axis and $z_{m,n}$ is the orthogonal morphing coordinate. More general hardware polytopes can be substituted without changing the EFIM derivation.

For the adopted orthogonal morphing architecture, the physical position of array element n can be written as

$$\mathbf{r}_{m,n} = \left[(n-1)\lambda/2 \quad z_{m,n} \right]^T. \quad (5)$$

Hence, for any two elements $i \neq j$,

$$\|\mathbf{r}_{m,i} - \mathbf{r}_{m,j}\|_2 = \sqrt{(i-j)^2 \lambda^2 / 4 + (z_{m,i} - z_{m,j})^2} \geq |i-j| \lambda / 2. \quad (6)$$

A minimum-spacing requirement $d_{\min,m} \leq \lambda/2$ is satisfied automatically by every feasible morphing configuration. For a stricter hardware spacing requirement, the explicit constraint $\|\mathbf{r}_{m,i} - \mathbf{r}_{m,j}\|_2 \geq d_{\min,m}$ is retained in the outer geometry-feasibility test; infeasible candidate shapes are rejected before solving the covariance subproblem. This additional hardware check does not alter the nuisance-aware EFIM derivation or the convexity of a fixed-shape covariance block.

3.1. Signal-consistent communication and probing

The Tx APs serve K_c single-antenna users. Let $\mathcal{C}_k \subseteq \{1, \dots, M_t\}$ denote the fixed user-centric serving cluster for user k , let $\chi_{t,k} = 1$ if $t \in \mathcal{C}_k$ and zero otherwise, and define the block selection matrix

$$\mathbf{D}_k^{\text{srv}} = \text{blkdiag}(\chi_{1,k} \mathbf{I}_{N_1}, \dots, \chi_{M_t,k} \mathbf{I}_{N_{M_t}}). \quad (7)$$

Full Cell-Free cooperation is the special case $\mathbf{D}_k^{\text{srv}} = \mathbf{I}$. We keep the cluster membership fixed in this paper in order to isolate nuisance-aware geometry design; AP-cluster selection can be placed in a slower outer layer. All reported numerical experiments use the full-cooperation specialization $\mathbf{D}_k^{\text{srv}} = \mathbf{I}$ so that the revised architectural clarification does not change the previously frozen numerical protocol. Stack the communication beamformers for user k into $\tilde{\mathbf{w}}_k$ and define $\mathbf{Q}_k = \tilde{\mathbf{w}}_k \tilde{\mathbf{w}}_k^H$, with the user-centric support constraint

$$\mathbf{Q}_k = \mathbf{D}_k^{\text{srv}} \mathbf{Q}_k \mathbf{D}_k^{\text{srv}}. \quad (8)$$

Tx AP t also emits a zero-mean probing signal with covariance $\mathbf{S}_t \succeq \mathbf{0}$. Probing sequences are orthogonal across Tx APs over an L -snapshot sensing frame and known at the central processor. Let $\mathbf{S} = \text{blkdiag}(\mathbf{S}_1, \dots, \mathbf{S}_{M_t})$. Treating probing as interference at the users gives

$$\text{SINR}_k = \frac{\tilde{\mathbf{h}}_k^H \mathbf{Q}_k \tilde{\mathbf{h}}_k}{\sum_{j \neq k} \tilde{\mathbf{h}}_k^H \mathbf{Q}_j \tilde{\mathbf{h}}_k + \tilde{\mathbf{h}}_k^H \mathbf{S} \tilde{\mathbf{h}}_k + \sigma_k^2}. \quad (9)$$

Communication and probing share each AP's power budget, and probing also contributes interference to users. Together with the user-centric serving constraints in Equation (8), these interactions couple communication and sensing across the cell-free network. The central processor jointly coordinates the covariances and fuses the separated bistatic returns.

After separation of Tx probing sequences, Rx AP r observes Tx AP t as

$$\mathbf{Y}_{tr} = \alpha_{tr} \mathbf{a}_r \mathbf{a}_t^H \mathbf{X}_t + \mathbf{N}_{tr}, \quad (10)$$

where $\alpha_{tr} \in \mathbb{C}$ is an unknown deterministic path coefficient, $\mathbf{X}_t \mathbf{X}_t^H = L \mathbf{S}_t$, and $\text{vec}(\mathbf{N}_{tr}) \sim \mathcal{CN}(\mathbf{0}, \zeta_r^2 \mathbf{I})$. Define

$$\mathbf{b}_t = \mathbf{X}_t^H \mathbf{a}_t, \quad \dot{\mathbf{b}}_t = \mathbf{X}_t^H \dot{\mathbf{a}}_t, \quad \mathbf{s}_{tr} = \mathbf{b}_t^* \otimes \mathbf{a}_r. \quad (11)$$

4. Nuisance-aware equivalent fisher information

4.1. Unknown complex gain elimination

For one separated link, let $\mathbf{D}_{tr} = [\dot{\mathbf{b}}_t^* \otimes \mathbf{a}_r, \mathbf{b}_t^* \otimes \dot{\mathbf{a}}_r]$. Treating the real and imaginary parts of α_{tr} as real nuisance parameters, Schur elimination gives

$$\mathbf{I}_{e,tr} = \frac{2|\alpha_{tr}|^2}{\zeta_r^2} \text{Re} \left\{ \mathbf{D}_{tr}^H \mathbf{\Pi}_{s_{tr}}^\perp \mathbf{D}_{tr} \right\}, \quad (12)$$

with $\mathbf{\Pi}_s^\perp = \mathbf{I} - \mathbf{s}\mathbf{s}^H / \|\mathbf{s}\|^2$. The projector removes every derivative component that is indistinguishable from a perturbation of the unknown complex gain.

Define the transmit illumination and centered tangent

$$q_t(\mathbf{S}_t) = \mathbf{a}_t^H \mathbf{S}_t \mathbf{a}_t, \quad (13)$$

$$\kappa_t(\mathbf{S}_t) = \dot{\mathbf{a}}_t^H \mathbf{S}_t \dot{\mathbf{a}}_t - \frac{|\mathbf{a}_t^H \mathbf{S}_t \dot{\mathbf{a}}_t|^2}{\mathbf{a}_t^H \mathbf{S}_t \mathbf{a}_t}, \quad q_t > 0, \quad (14)$$

For comparison, if the complex bistatic gain were treated as known (or omitted from the parameter vector), the transmit-angle information would use the unprojected quantity

$$d_t^{\text{KG}}(\mathbf{S}_t) = \dot{\mathbf{a}}_t^H \mathbf{S}_t \dot{\mathbf{a}}_t. \quad (15)$$

where the superscript KG denotes the known-gain surrogate. Since the second term in Equation (14) is nonnegative, $0 \leq \kappa_t \leq d_t^{\text{KG}}$. We define the nuisance-retention ratio and known-gain optimism factor

$$\eta_t = \frac{\kappa_t}{d_t^{\text{KG}}} \in [0, 1], \quad \mathcal{O}_t = \frac{d_t^{\text{KG}}}{\kappa_t}, \quad (16)$$

with $\mathcal{O}_t = \infty$ when $d_t^{\text{KG}} > 0$ but $\kappa_t = 0$. Hence nuisance omission is not equivalent to a fixed dB offset: it can change whether a direction is identifiable at all.

And the receive centered aperture

$$\mathbf{v}_r = \dot{\mathbf{a}}_r^H (\mathbf{I} - \mathbf{a}_r \mathbf{a}_r^H) \dot{\mathbf{a}}_r. \quad (17)$$

Then the separated-link EFIM reduces to

$$\mathbf{I}_{e,tr} = \frac{2L|\alpha_{tr}|^2}{\zeta_r^2} \begin{bmatrix} \kappa_t & 0 \\ 0 & q_t \mathbf{v}_r \end{bmatrix}. \quad (18)$$

The cross term between the angle of departure (AoD) and angle of arrival (AoA) vanishes for separated rank-one Kronecker links with independent unknown gains. This result does not generally hold for unresolved returns.

4.2. Exact probing identifiability

Proposition 1 (probing identifiability): Let $q_t > 0$, $\mathbf{u} = \mathbf{S}_t^{1/2} \dot{\mathbf{a}}_t$, and $\mathbf{v} = \mathbf{S}_t^{1/2} \mathbf{a}_t$. Then

$$\kappa_t = \|\mathbf{\Pi}_v^\perp \mathbf{u}\|^2, \quad (19)$$

so

$$\kappa_t > 0 \iff \Pi_{S_t^{1/2} \mathbf{a}_t}^\perp S_t^{1/2} \dot{\mathbf{a}}_t \neq \mathbf{0}. \quad (20)$$

Every rank-one S_t yields $\kappa_t = 0$; hence $\text{rank}(S_t) \geq 2$ is necessary but not sufficient for nonzero AoD information.

Proof: Equation (14) is $\|\mathbf{u}\|^2 - |\mathbf{v}^H \mathbf{u}|^2 / \|\mathbf{v}\|^2$, the squared norm of the component of $\mathbf{u}$ orthogonal to $\mathbf{v}$. A rank-one square root maps both vectors into the same one-dimensional range. Higher rank does not by itself prevent collinearity.

Table 2 summarizes the qualitative information gap created by omitting the unknown complex gain.

Table 2. Why nuisance omission can qualitatively overstate geometry information.

Model	Tx-Angle Metric	Rank-One Probing Consequence
Known-gain/raw surrogate	$d_t^{\text{KG}} = \dot{\mathbf{a}}_t^H S_t \dot{\mathbf{a}}_t$	Can remain positive
Unknown-gain EFIM	$\kappa_t = d_t^{\text{KG}} - \mathbf{a}_t^H S_t \dot{\mathbf{a}}_t ^2 / q_t$	Exactly zero

In a single-Rx complementary geometry, receive-angle information contributes only one position direction. A rank-one probing design with $d_t^{\text{KG}} > 0$ may therefore appear localizable under the known-gain surrogate, while Proposition 1 removes all Tx-angle information and leaves a rank-deficient two-dimensional nuisance-aware position EFIM. The associated CRB optimism is unbounded rather than a small numerical mismatch.

For isotropic probing $S_t = (P_t^s/N_t)\mathbf{I}$, $q_t = P_t^s/N_t$ and $\kappa_t = (P_t^s/N_t)v_t$.

4.3. Identifiable morphing aperture

Using $\dot{\mathbf{a}}_m = j\mathbf{D}_m \mathbf{a}_m$ gives

$$v_m = \frac{1}{N_m} \sum_{n=1}^{N_m} (d_{m,n} - \bar{d}_m)^2, \quad \bar{d}_m = \frac{1}{N_m} \sum_n d_{m,n}. \quad (21)$$

A common phase-rate shift $\mathbf{d}_m \mapsto \mathbf{d}_m + c\mathbf{1}$ is annihilated by $\mathbf{I} - \mathbf{a}_m \mathbf{a}_m^H$. Thus only differential phase-rate changes are identifiable in the adopted far-field unknown-gain model. For fixed θ_m , v_m is convex quadratic in the element displacements, which are coupled through the centering term. Over a compact box or polytope, at least one maximizer lies at a vertex.

4.4. Nonorthogonal or unresolved probing

If link returns cannot be separated, stack the mean as $\boldsymbol{\mu}(\boldsymbol{\theta}, \boldsymbol{\alpha}) = \mathbf{B}(\boldsymbol{\theta})\boldsymbol{\alpha}$ with noise covariance $\mathbf{C}$. The controlling expression is

$$I_e(\boldsymbol{\theta}) = 2 \text{Re} \left\{ \mathbf{D}_\theta^H \mathbf{C}^{-1/2} \Pi_{\mathbf{C}^{-1/2} \mathbf{B}}^\perp \mathbf{C}^{-1/2} \mathbf{D}_\theta \right\}. \quad (22)$$

Cross-AP and AoD–AoA terms generally remain in this case.

5. Calibration sensitivity and hybrid EFIM

The nominal branch assumes known post-calibration geometry, whereas the revised calibration-aware branch in Section 7.4. places residual calibration uncertainty directly inside the geometry/covariance design loop. This section uses the same complete-information principle that underlies hybrid-CRB and array-shape self-calibration analyses [20,25,26], but applies it to residual geometry errors around a commanded ISAC aperture. Let ψ_m be the nominal orientation, δ_m a residual orientation bias, and e_m an execution error. For global target bearing $\theta_m(\mathbf{p})$,

$$\vartheta_m = \theta_m(\mathbf{p}) - \psi_m - \delta_m, \quad \mathbf{z}_m^{\text{true}} = \mathbf{z}_m + e_m. \quad (23)$$

The realized phase rate is

$$d_{m,n} = c_{m,n} \cos \vartheta_m - k_0(z_{m,n} + e_{m,n}) \sin \vartheta_m, \quad (24)$$

where $c_{m,n} = (n-1)\pi$ and $k_0 = 2\pi/\lambda$.

For a fixed calibration state, unknown complex gains are first eliminated with the same projector construction. If γ collects calibration variables, the resulting prior-regularized local information matrix is

$$\mathbf{J}_{\text{loc}} = \begin{bmatrix} \mathbf{J}_{pp} & \mathbf{J}_{p\gamma} \\ \mathbf{J}_{\gamma p} & \mathbf{J}_{\gamma\gamma} + \mathbf{\Lambda}_\gamma \end{bmatrix}, \quad (25)$$

with position EFIM

$$\mathbf{J}_p^{\text{loc}} = \mathbf{J}_{pp} - \mathbf{J}_{p\gamma}(\mathbf{J}_{\gamma\gamma} + \mathbf{\Lambda}_\gamma)^{-1} \mathbf{J}_{\gamma p}. \quad (26)$$

For random calibration errors, the exact hybrid bound requires averaging the complete data FIM before taking the Schur complement.

$$\mathbf{J}_H^{\text{full}} = \mathbb{E}_\gamma[\mathbf{J}_D^{\text{full}}] + \mathbf{J}_{\text{prior}}, \quad (27)$$

and only then is the full nuisance block eliminated. In general,

$$\text{Schur}(\mathbb{E}[\mathbf{J}]) \neq \mathbb{E}[\text{Schur}(\mathbf{J})]. \quad (28)$$

Scenario-wise local CRBs are retained only as deployment-risk summaries.

The exact hybrid CRB and CVaR₉₅ measure different aspects of estimation performance and need not coincide. We distinguish them as follows:

$$\begin{aligned} C_H &\triangleq \text{tr} \left(\left[\text{Schur}_{p|\gamma} \left(\mathbb{E}_\gamma[\mathbf{J}_D^{\text{full}}] + \mathbf{J}_{\text{prior}} \right) \right]^{-1} \right), \\ c_s &\triangleq \text{tr} \left(\left[\text{Schur}_{p|\gamma} \left(\mathbf{J}_D^{\text{full}}(\gamma_s) + \mathbf{J}_{\text{prior}} \right) \right]^{-1} \right). \end{aligned} \quad (29)$$

The calibration-aware design criterion is the hybrid quantity C_H , for which the complete information is averaged before nuisance elimination. For fixed geometry and covariances, each c_s is the conditional/local CRB for scenario s . We report their mean, quantiles, and CVaR₉₅ to assess deployment risk. Since Schur complementation and matrix inversion are nonlinear, C_H generally differs from $\mathbb{E}[c_s]$. We assume neither equality nor a fixed ordering between C_H and $\text{CVaR}_{95}(\{c_s\})$.

Proposition 2 (common-displacement invariance): Replacing e_m by $e_m + c\mathbf{1}$ leaves the gain-eliminated EFIM unchanged. Nuisance elimination removes the common phase and phase-rate components. Only relative actuator displacements can affect identifiable angular information.

Proposition 3 (orientation-information saturation): If a scalar local bearing has data information I_m and an independent orientation bias $\delta_m \sim \mathcal{N}(0, \sigma_{\delta,m}^2)$ enters only through $\theta_m - \delta_m$, then

$$I_{m,\text{eff}} = \frac{I_m}{1 + I_m \sigma_{\delta,m}^2}, \quad I_{m,\text{eff}} \leq \sigma_{\delta,m}^{-2}. \quad (30)$$

Proof: The data-FIM block for (θ_m, δ_m) is $I_m \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix}$. Adding prior information $\sigma_{\delta,m}^{-2}$ to the nuisance block and taking the Schur complement gives Equation (30).

The scalar ceiling in Equation (30) is a special case of the full hybrid-Schur expression. If the complete prior-regularized information matrix for position and a calibration nuisance vector δ is

$$\mathbf{J}_{\text{full}} = \begin{bmatrix} \mathbf{J}_{pp} & \mathbf{J}_{p\delta} \\ \mathbf{J}_{\delta p} & \mathbf{J}_{\delta\delta} + \mathbf{J}_{\text{prior}} \end{bmatrix}, \quad (31)$$

then the exact hybrid position EFIM is

$$\mathbf{J}_p^{\text{hyb}} = \mathbf{J}_{pp} - \mathbf{J}_{p\delta}(\mathbf{J}_{\delta\delta} + \mathbf{J}_{\text{prior}})^{-1} \mathbf{J}_{\delta p}. \quad (32)$$

If a position direction is coupled only through a poorly known calibration mode, its information remains limited by that mode's prior precision, even as the data SNR increases. This generalizes the scalar information ceiling in Equation (30).

An analogous local ceiling follows for independent element execution errors. Let $\mathbf{P}_m^\perp = \mathbf{I} - \mathbf{1}\mathbf{1}^T/N_m$, $\mathbf{v}_m = \mathbf{d}_m^T \mathbf{P}_m^\perp \mathbf{d}_m/N_m$, and write the gain-eliminated bearing information as $I_{\theta,m} = \beta_m \mathbf{v}_m$. If $e_{m,n} \sim \mathcal{N}(0, \sigma_{e,m}^2)$ independently and no other geometry nuisance is present, eliminating the differential element errors gives

$$I_{\theta,m,\text{eff}} = \frac{\beta_m \mathbf{v}_m}{1 + \beta_m k_0^2 \cos^2 \vartheta_m \sigma_{e,m}^2 / N_m}. \quad (33)$$

The common element-displacement mode is absent from Equation (33), consistent with Proposition 2. Correlated actuator errors do not reduce to a single scalar denominator and must instead be handled with the full matrix Schur complement and a physically justified covariance model.

Figure 2 shows the frozen orientation negative control and the residual tail penalty after covariance reallocation. In the frozen orientation audit, treating orientation as unknown raises CVaR₉₅ of the scenario-wise local position CRBs from 8.733 to 18.354 m², a 3.23-dB penalty. Calibration-aware covariance reallocation gives 18.386 m², providing no improvement (a gain of −0.007 dB). These are tail-risk statistics of local CRBs, not the exact hybrid CRB in Equation (32). This experiment serves as a negative control.




Figure 2. Frozen 128-scenario orientation negative control at $\sigma_\delta = 2^\circ$. Values are CVaR₉₅ summaries of scenario-wise conditional/local position CRBs, not exact hybrid CRBs. The 3.23-dB tail penalty is not recovered by covariance reallocation; the calibration-aware joint result corresponds to a -0.007 -dB change under the nuisance model.

6. Position EFIM and angular diversity

For an AP at $\mathbf{p}_m = [x_m, y_m]^T$,

$$\mathbf{g}_m = \frac{\partial \theta_m}{\partial \mathbf{p}} = \frac{1}{r_m} \begin{bmatrix} -\sin \theta_m \\ \cos \theta_m \end{bmatrix}, \quad r_m = \|\mathbf{p} - \mathbf{p}_m\|. \quad (34)$$

Under link-separable probing,

$$\mathbf{I}_p = \sum_m \mathbf{I}_m \mathbf{g}_m \mathbf{g}_m^T. \quad (35)$$

Writing $w_m = \mathbf{I}_m / r_m^2$ gives the determinant identity

$$\det \mathbf{I}_p = \sum_{m < m'} w_m w_{m'} \sin^2(\theta_m - \theta_{m'}). \quad (36)$$

Hence a finite two-dimensional CRB requires at least two positive-weight, nonparallel directions. On this domain,

$$\text{CRB}(\mathbf{p}) = \text{tr}(\mathbf{I}_p^{-1}) = \frac{\sum_m w_m}{\sum_{m < m'} w_m w_{m'} \sin^2(\theta_m - \theta_{m'})}. \quad (37)$$

This separates local aperture sensitivity from deployment angular diversity.

7. Joint covariance and shape design

7.1. Problem formulation

Let $\mathbf{E}_t$ select Tx AP t from a stacked covariance. The nominal problem is

$$\min_{\{\mathbf{Q}_k\}, \{\mathbf{S}_t\}, \{z_m\}} \text{tr}(\mathbf{I}_p^{-1}) \quad (38a)$$

$$\text{s.t.} \quad \text{SINR}_k \geq \Gamma_k, \quad \forall k, \quad (38b)$$

$$\sum_k \text{tr}(\mathbf{E}_t \mathbf{Q}_k) + \text{tr}(\mathbf{S}_t) \leq P_t, \quad \forall t, \quad (38c)$$

$$\mathbf{Q}_k \succeq \mathbf{0}, \quad \mathbf{S}_t \succeq \mathbf{0}, \quad (38d)$$

$$\mathbf{Q}_k = \mathbf{D}_k^{\text{srv}} \mathbf{Q}_k \mathbf{D}_k^{\text{srv}}, \quad \forall k, \quad (38e)$$

$$q_t(\mathbf{S}_t) \geq \underline{q}_t > 0, \quad \forall t, \quad (38f)$$

$$|z_{m,n}| \leq \tilde{z}_m, \quad (38g)$$

$$\|\mathbf{r}_{m,i} - \mathbf{r}_{m,j}\|_2 \geq d_{\min,m}, \quad \forall m, i \neq j. \quad (38h)$$

The illumination floor keeps the unknown-gain nuisance block regular and prevents a conic-closure solution from driving the nominal echo to zero while retaining a large derivative. We use $\underline{q}_t = 0.05P_t$ as a conservative numerical regularity floor, not as a tuned sensing-power allocation. The constraint is strongly inactive at the reported nominal optimum, where the minimum achieved illumination is $0.766P_t$, and it remains nonbinding in the revised 2° calibration-aware runs, where the minimum training-scenario illumination ratios are 0.427–0.663. Equation (6) further shows that (38h) is automatically satisfied for $d_{\min,m} \leq \lambda/2$ under the adopted orthogonal morphing architecture; stricter spacing requirements are enforced by rejecting infeasible outer shape candidates. Section 8.3. additionally stress-tests $d_{\min}/\lambda \in \{0.5, 0.6, 0.75\}$ and verifies zero spacing violation under the explicit pre-SDP feasibility gate.

7.2. Fixed-shape SDP and regularity

For fixed shapes, introduce $\tau_t \leq \kappa_t(\mathbf{S}_t)$. Under $q_t > 0$,

$$\begin{bmatrix} \dot{\mathbf{a}}_t^H \mathbf{S}_t \dot{\mathbf{a}}_t - \tau_t & \dot{\mathbf{a}}_t^H \mathbf{S}_t \mathbf{a}_t \\ \mathbf{a}_t^H \mathbf{S}_t \dot{\mathbf{a}}_t & \mathbf{a}_t^H \mathbf{S}_t \mathbf{a}_t \end{bmatrix} \succeq \mathbf{0} \quad (39)$$

is exactly equivalent to $\tau_t \leq \kappa_t$. On the finite-localization domain $\mathbf{I}_p \succ \mathbf{0}$,

$$\begin{bmatrix} \mathbf{Z} & \mathbf{I}_2 \\ \mathbf{I}_2 & \mathbf{I}_p \end{bmatrix} \succeq \mathbf{0} \iff \mathbf{Z} \succeq \mathbf{I}_p^{-1}. \quad (40)$$

Thus minimizing $\text{tr}(\mathbf{Z})$ gives the A-optimal objective; these are standard Schur-complement semidefinite representations [28]. The user-centric support relations Equation (38e) are linear equality constraints and therefore preserve the fixed-shape SDP structure. If $\mathbf{I}_p$ is singular, this block linear matrix inequality (LMI) is infeasible rather than a finite CRB representation.

7.3. Exact communication rank-one recovery

Proposition 4: Let $\mathbf{Q}_k^*$ be a feasible relaxed communication covariance and $g_k = \tilde{\mathbf{h}}_k^H \mathbf{Q}_k^* \tilde{\mathbf{h}}_k > 0$. Define

$$\mathbf{w}_k^* = \frac{\mathbf{Q}_k^* \tilde{\mathbf{h}}_k}{\sqrt{g_k}}, \quad \hat{\mathbf{Q}}_k = \mathbf{w}_k^* (\mathbf{w}_k^*)^H. \quad (41)$$

Then $\mathbf{0} \preceq \hat{\mathbf{Q}}_k \preceq \mathbf{Q}_k^*$, user k 's desired received power is preserved, cross-user interference and selected per-AP communication power cannot increase, and the probing-only EFIM remains unchanged because $\{\mathbf{S}_t\}$ is unchanged. Moreover, if $\mathbf{Q}_k^* = \mathbf{D}_k^{\text{srv}} \mathbf{Q}_k^* \mathbf{D}_k^{\text{srv}}$, then $\mathbf{w}_k^*$ lies in the range of $\mathbf{D}_k^{\text{srv}}$ and the recovered $\hat{\mathbf{Q}}_k$ satisfies the same user-centric support constraint. Hence the communication rank relaxation is exact for the adopted dedicated-probing sensing model.

Proof: Generalized Cauchy–Schwarz in the seminorm induced by $\mathbf{Q}_k^*$ implies $|\mathbf{x}^H \mathbf{Q}_k^* \tilde{\mathbf{h}}_k|^2 \leq (\mathbf{x}^H \mathbf{Q}_k^* \mathbf{x}) g_k$ for every $\mathbf{x}$, which is equivalent to $\hat{\mathbf{Q}}_k \preceq \mathbf{Q}_k^*$. Direct substitution preserves the desired quadratic form; all other positive-semidefinite (PSD) quadratic forms and selected block traces are nonincreasing.

7.4. Calibration-aware joint geometry and covariance design

To incorporate calibration errors into the geometry optimization itself, let $\mathcal{G}_{N_{\text{sc}}} = \{\gamma_s\}_{s=1}^{N_{\text{sc}}}$ be a fixed common-random-number bank of orientation and actuator-error realizations. For a commanded shape $\mathbf{z}$, the execution-error component of γ_s perturbs the realized steering geometry, while the residual orientation bias is retained as a prior-regularized nuisance coordinate. Let $\tilde{\mathbf{J}}_s(\mathbf{z}, \{\mathbf{S}_t\})$ denote the resulting real position–orientation information matrix after the unknown complex bistatic gains have been eliminated exactly by the projector construction of Section 4. The tractable sample-average hybrid design criterion is

$$\bar{\mathbf{H}}(\mathbf{z}, \{\mathbf{S}_t\}) = \frac{1}{N_{\text{sc}}} \sum_{s=1}^{N_{\text{sc}}} \tilde{\mathbf{J}}_s + \mathbf{J}_{\text{prior}}. \quad (42)$$

Let $\mathbf{E}_p = [\mathbf{I}_2, \mathbf{0}]$ select the target-position coordinates. For fixed commanded geometry, the calibration-aware covariance subproblem is

$$\min_{\{\mathbf{Q}_k\}, \{\mathbf{S}_t\}, \mathbf{Z}_H} \text{tr}(\mathbf{Z}_H) \quad (43a)$$

$$\text{s.t.} \quad \begin{bmatrix} \mathbf{Z}_H & \mathbf{E}_p \\ \mathbf{E}_p^T & \bar{\mathbf{H}} \end{bmatrix} \succeq \mathbf{0}, \quad (43b)$$

$$\text{SINR}_k \geq \Gamma_k, \quad \forall k, \quad (43c)$$

$$\sum_k \text{tr}(\mathbf{E}_t \mathbf{Q}_k) + \text{tr}(\mathbf{S}_t) \leq P_t, \quad \forall t, \quad (43d)$$

$$\mathbf{Q}_k \succeq \mathbf{0}, \quad \mathbf{S}_t \succeq \mathbf{0}, \quad (43e)$$

$$\mathbf{Q}_k = \mathbf{D}_k^{\text{srv}} \mathbf{Q}_k \mathbf{D}_k^{\text{srv}}, \quad \forall k, \quad (43f)$$

$$\mathbf{a}_{t,s}^H \mathbf{S}_t \mathbf{a}_{t,s} \geq \underline{q}_t, \quad \forall t, s. \quad (43g)$$

The scenario-wise illumination constraints preserve the regularity required for exact complex-gain elimination throughout the sampled calibration set. The averaging in Equation (42) is performed before eliminating the retained orientation nuisance. Because element execution errors are sampled deployment states, this criterion is referred to as a hybrid SAA design criterion; the complete-FIM exact hybrid CRB

of Section 5. remains the theoretical reference.

Proposition 5 (fixed-shape calibration-aware convexity): For fixed commanded shapes and a fixed calibration-scenario bank, Equation (43) is an SDP. The deterministic communication rank-one recovery of Proposition 4 remains valid.

Proof: For fixed commanded geometry, all scenario steering vectors and derivatives are constants. The exact gain-eliminated link information is represented by the same Schur-LMI epigraph used in the nominal block, so the retained position–orientation information is affine in the probing covariances and auxiliary tangent-information variables; hence the fixed-shape formulation remains semidefinite representable. The position-CRB epigraph is a Schur LMI, while the SINR, per-AP power, scenario-illumination, PSD, and fixed user-centric support constraints are convex/affine. Replacing the relaxed communication covariances by the dominated rank-one matrices of Proposition 4 leaves $\{\mathbf{S}_t\}$, and the sensing information, unchanged.

Define the optimized hybrid value function

$$F_H(\mathbf{z}) = \min \text{tr}(\mathbf{Z}_H) \quad (44)$$

subject to Equation (43) for fixed $\mathbf{z}$.

Starting from the nominal optimized commanded shape, both admissible endpoints of every Tx and Rx displacement coordinate are evaluated using the same training scenario bank. Unlike the nominal receive-coordinate shortcut, a calibration-aware Rx update is also re-solved because orientation and actuator errors modify the scenario-dependent steering geometry. Candidate shapes that violate Equation (38h) are rejected before the SDP call. Once the training optimization is complete, geometry and covariance are frozen before an independent out-of-sample calibration bank is generated. The hybrid-SAA CRB is then reported together with the mean, 95th percentile, and CVaR₉₅ summaries of scenario-wise conditional/local CRBs, preserving the distinction stated in Section 5.

7.5. Safeguarded shape updates

Let $F(\mathbf{z})$ denote the optimized fixed-shape objective: the nominal A-optimal CRB for the nominal branch and $F_H(\mathbf{z})$ in Equation (44) for the calibration-aware branch. The revised outer search uses a deterministic endpoint coordinate sweep together with the strict acceptance rule

$$F(\mathbf{z}') \leq F(\mathbf{z}) - \epsilon_{\text{acc}}, \quad \epsilon_{\text{acc}} = \max\{\epsilon_{\text{abs}}, \epsilon_{\text{rel}}|F(\mathbf{z})|\}. \quad (45)$$

Starting from any feasible commanded shape, one nominal sweep performs the following steps.

- (1) Solve the fixed-shape covariance SDP and record its optimized objective as the incumbent value.
- (2) For each receive coordinate, use $\mathbf{v}_r$ only as an analytical prescreen of the two endpoints. Because a receive-only update leaves the communication constraints and probing covariances unchanged, its CRB can be evaluated under the incumbent covariances without a new conic solve. The move is accepted only when this feasible incumbent-covariance value satisfies Equation (45); consequently the newly optimized fixed-shape value is also strictly lower.
- (3) For each transmit coordinate, test the two endpoints one at a time. For each distinct feasible candidate, including the spacing test Equation (38h), rebuild the steering quantities, resolve the fixed-shape SDP, and accept only if Equation (45) holds.

- (4) After the sweep, apply the deterministic communication rank-one recovery of Proposition 4; the sensing covariances and objective are unchanged.

For the calibration-aware branch, every Tx and Rx endpoint is instead evaluated by resolving Equation (43), because the scenario-dependent calibration geometry removes the nominal receive-only shortcut.

Proposition 6 (finite termination and endpoint stability): Assume $\epsilon_{\text{abs}} > 0$ and $\epsilon_{\text{rel}} \geq 0$ in Equation (45). Under deterministic coordinate ordering and the finite endpoint candidate set, both safeguarded searches terminate after finitely many accepted moves. At termination no admissible one-coordinate endpoint replacement decreases the corresponding optimized objective by at least ϵ_{acc} .

Proof: Every accepted update is feasible and decreases the nonnegative objective by at least $\epsilon_{\text{acc}} \geq \epsilon_{\text{abs}} > 0$. Hence, if N_{acc} moves are accepted from the initial value $F(\mathbf{z}^{(0)})$, then

$$N_{\text{acc}} \leq \left\lfloor \frac{F(\mathbf{z}^{(0)})}{\epsilon_{\text{abs}}} \right\rfloor, \quad (46)$$

so infinitely many accepted moves are impossible. Independently, after a coordinate is visited it can occupy only one of its registered endpoints, so the reachable endpoint state set is finite and an accepted-state cycle is excluded by strict decrease. Termination after a complete sweep with no accepted coordinate is exactly the stated coordinatewise endpoint-stability condition.

This certificate is deliberately finite-set: it rules out further accepted one-coordinate endpoint moves under the registered search and threshold, but it does not imply stationarity with respect to interior continuous perturbations, continuous-domain KKT stationarity, or global optimality of the underlying mixed nonconvex problem.

Computational complexity: Let $N_{\text{Tx}} = \sum_t N_t$, $N_{\text{Rx}} = \sum_r N_r$, and let $d_H = 2 + d_\gamma$ be the dimension of the retained real hybrid design vector, where $d_\gamma = M_t$ for the main orientation-nuisance branch. The complex bistatic gains have already been eliminated exactly, while sampled actuator errors perturb the scenario geometry and therefore do not enlarge d_H in this main solver. The covariance variables contain

$$n_{\text{cov}} = K_c N_{\text{Tx}}^2 + \sum_t N_t^2 \quad (47)$$

real degrees of freedom up to lower-order epigraph variables. A nominal fixed-shape SDP contains the communication/probing PSD cones and the 2×2 localization Schur LMI; the calibration-aware problem additionally contains a $(d_H + 2) \times (d_H + 2)$ hybrid Schur LMI and $N_{\text{sc}} M_t$ scalar illumination constraints. A conservative dense interior-point proxy is cubic in the KKT system dimension per Newton step, while practical conic solvers exploit the separate PSD/LMI blocks; we report both solve counts and measured wall-clock times. After the incumbent solve, a nominal sweep requires at most $2N_{\text{Tx}}$ additional SDP resolves because Rx candidates use the analytical prescreen/evaluation above, whereas a calibration-aware sweep requires at most $2(N_{\text{Tx}} + N_{\text{Rx}})$ robust SDP resolves. Repeated endpoints, spacing rejections, and early termination reduce these upper bounds. These are per-sweep outer-solve counts; the dominant cost in the calibration-aware branch remains the scenario-expanded fixed-shape SDP, not the endpoint bookkeeping itself. Because interior-point implementations exploit different cone factorizations, we do not present the dense cubic proxy as a hardware-independent runtime law; instead, Section 8. reports

measured solve counts, solver time, wall time, and memory for the actual implementation used.

8. Numerical evaluation

All sensing bounds in this section use the nuisance-aware EFIM derived above; known-gain Fisher information is never reused. We organize the evidence in the same order as the claims: first internal correctness and algorithmic behavior, then the physical 28-GHz benchmark, followed by the journal-scale generalization chain of 100 random full-envelope geometries, 960 paired shape–trial comparisons, 121 independently reoptimized target locations, and 2560 finite-sample ML cases, and finally the calibration boundary. Except in the ML subsection, performance is reported as position CRB in m^2 . Table 3 summarizes the validation layers and their principal acceptance outcomes.

Table 3. Extended validation matrix beyond the nominal design point.

Axis	Scale	Main Outcome
Random geometry	100 shapes	Joint wins 100/100; median gain 3.07 dB
Paired physical out-of-sample	$20 \times 48 = 960$	Joint wins 960/960; median shape-mean 2.65 dB
Spatial redesign	$11 \times 11 = 121$	Joint wins 121/121; median gain 5.68 dB
Finite-sample ML	$4 \times 5 \times 128 = 2560$	0 failures; high-SNR gap to RMS-PEB < 0.45 dB
Orientation sweep	7 uncertainty levels	CVaR ₉₅ rises $8.08 \rightarrow 71.67 \text{ m}^2$; robust change ≤ 0.051 dB
Movable/fluid-antenna algorithm baselines	5 SINR targets + 20 initializations + 3 envelopes	lowest CRB at all sweep points; 20/20 initialization wins
Link-level communication	5 SINR targets, 10^6 quadrature phase-shift keying symbols/user/point	bit error rate $0.2391 \rightarrow 2.383 \times 10^{-3}$; all consistency gates pass
Illumination-floor sensitivity	4 floors $\times$ 2 SINR loads	inactive at 3 dB; active but only 1.48% CRB spread at 9 dB
Minimum-spacing stress	$d_{\min}/\lambda = \{0.5, 0.6, 0.75\}$	automatic at $\leq 0.5\lambda$; stricter cases satisfy the explicit feasibility gate
Computational profile	nominal + 3 robust regimes	39 SDP solves/7.54 s nominal; 26 robust candidates/352–623 s aggregate solver time
Convergence certificate	finite endpoint set + strict decrease	finite accepted-move bound; endpoint stability only, not continuous KKT/global optimality
Hybrid vs. conditional risk	hybrid-SAA + frozen out-of-sample scenarios	average complete FIM before Schur; CVaR ₉₅ only over scenario-wise local CRBs
Calibration-aware redesign	3 separated regimes + 29-D audit	0/26 endpoint moves in each regime; frozen 29-D $D1 = D2$

8.1. Internal correctness and algorithmic behavior

For 1000 random feasible shapes, probing matrices, and reflection coefficients, we compare the complete real-parameter FIM, its Schur complement, the projector EFIM in Equation (12), and the separated-link closed form. The maximum relative Frobenius discrepancies are 4.91×10^{-15} and 5.66×10^{-15} , while the largest residual AoD–AoA entry after nuisance elimination is 9.06×10^{-13} . A second 1000-trial audit restricted to rank-one probing gives a largest absolute transmit centered-information residual of 5.08×10^{-13} . These machine-precision checks validate both the nuisance projection and Proposition 1.

Equation (37) also exposes a purely geometric failure mode. For two equal-information observations, the CRB penalty relative to orthogonal bearings is $1/\sin^2(\Delta\theta)$; Figure 3 shows its divergence as the directions become parallel.




Figure 3. Angular-diversity penalty for two equal-information observations, normalized to orthogonal bearings. The penalty increases sharply as the two bearings approach collinearity and vanishes toward orthogonal viewing.

The normalized algorithmic reference uses three dual-functional APs on a 75-m ring, two users, 12 Tx and 12 Rx elements in total, normalized total power 3, $L = 64$, $\tilde{z} = 0.5\lambda$, $\sigma_k^2 = 10^{-2}$, a common 3-dB SINR target, and $q_t = 0.05P_t$. Fixed-shape SDPs are solved with CLARABEL to 10^{-7} relative and absolute tolerances. Starting from zero displacements, one safeguarded shape sweep reduces the CRB from 5.048 to 1.2893 m²; a second sweep accepts no coordinate change. A random-vertex initialization reaches 1.2890 m² after its first sweep and is likewise unchanged by the second. Figure 4 shows the two-start convergence behavior and the endpoint-stable stopping point.

The production calibration-aware searches provide a separate empirical check of the same stopping event: in the orientation-only, actuator-only, and mixed-error regimes, each complete Tx/Rx endpoint sweep evaluates all 26 admissible candidates, accepts zero moves, and reports the endpoint-stability flag as true. These observations verify termination for the reported runs but are not used to strengthen Proposition 6 beyond its finite endpoint domain.




Figure 4. Normalized safeguarded shape audit. The revised strict acceptance rule guarantees finite termination at a coordinatewise endpoint-stable geometry; this finite-set certificate does not imply continuous-domain KKT stationarity or global optimality.

The deterministic recovery of Proposition 4 leaves the CRB unchanged; the largest SINR, per-AP power, and covariance-domination residuals are 3.1×10^{-7} , 7.7×10^{-8} , and 1.9×10^{-8} . The optimized probing covariances retain a second eigenmode carrying 5.92–8.06% of sensing power, consistent with the rank-one no-AoD result.

8.2. Algorithm-level movable/fluid-antenna baselines

To separate the proposed geometry-search rule from the nuisance-aware objective itself, we add algorithm-level comparisons against three representative movable/fluid-antenna ISAC optimization strategies. We implement *common-system adaptations*, rather than comparing numerical values across incompatible published systems: **(i)** A Qin-style AO/SCA geometry block motivated by the CRB-minimization framework in [29]; **(ii)** The MM/Rosen projected-gradient geometry-update philosophy of Chen *et al.* [11]; and **(iii)** The second-order-Taylor AO geometry block used for multi-static fluid-antenna AP-position design in Lou *et al.* [13]. For every method, the sensing metric is the same nuisance-aware position CRB, and every accepted geometry is re-evaluated by the same exact fixed-shape covariance SDP with the same SINR, per-AP power, illumination, displacement-envelope, and spacing constraints. The comparison therefore changes the geometry-update strategy, not the resource model or the sensing bound.

Figure 5a varies the common user SINR target from -3 to 9 dB. The proposed endpoint method attains the lowest CRB at all five operating points: 1.143, 1.195, 1.289, 1.463, and 1.699 m^2 , respectively. Relative to the best of the three literature-adapted optimizers at each SINR, the reductions are 0.73%, 0.52%, 0.57%, 1.47%, and 2.70%. The fixed-array control increases from 4.027 to 6.896 m^2 over the same sweep. Thus the expected communication–sensing tradeoff remains visible while the proposed method remains competitive throughout the tested load range. This dependence on the communication SINR target is also a Cell-Free-ISAC specificity check: the sensing optimum changes because the same distributed AP budgets support clustered user transmission and dedicated probing; in the sensing-only

multistatic specialization ($K_c = 0$), this SINR-driven coupling vanishes.

Initialization sensitivity is tested independently using 20 common random feasible geometries at 3 dB. As shown in Figure 5b, the proposed method wins all 20 paired trials. Its median CRB is 1.2883 m² with an interquartile range (IQR) of only 0.0037 m², compared with medians of 1.4585, 1.5462, and 1.4598 m² and IQRs of 0.1092, 0.3036, and 0.1611 m² for the Qin-, Chen-, and Lou-style updates, respectively. The proposed worst case (1.2951 m²) remains below the best value attained by each individual literature-adapted baseline (1.3135, 1.3068, and 1.3175 m²). Relative to the best literature-adapted result on each seed, the proposed reduction ranges from 1.96% to 17.54%, with a median of 10.60%. This experiment indicates that the main practical advantage of the safeguarded finite-set search is not an exaggerated single-initialization gain, but reduced sensitivity to poor local geometric solutions.




Figure 5. Common-system comparison with representative movable/fluid-antenna ISAC geometry-update strategies. **(a)** Across the common SINR sweep, the safeguarded endpoint design attains the lowest position CRB at all five operating points; **(b)** Over 20 common random feasible initializations at 3 dB, the proposed method exhibits a markedly tighter CRB distribution and wins every paired trial, indicating substantially lower sensitivity to poor local geometry solutions.

The aperture-envelope sweep provides a separate geometry-control check. For $\tilde{z}/\lambda \in \{0.25, 0.50, 0.75\}$, the proposed CRBs are 2.3194, 1.2893, and 0.8543 m², respectively, corresponding to a 63.2% reduction from the smallest to the largest envelope. The proposed method is again lowest at all three points; its reductions relative to the best literature-adapted method are 0.51%, 0.57%, and 1.16%. Hence the ranking is not confined to one chosen mobility range.

The stronger robustness is purchased with additional exact evaluations. At the 3-dB zero initialization, the proposed method uses 39 fixed-shape SDP solves, versus 28, 40, and 31 for the adapted Qin, Chen, and Lou methods, respectively; its 7.54-s wall time lies between the lighter Qin update (6.49 s) and the second-order-Taylor Lou update (19.01 s). We therefore interpret the endpoint search as a slow-timescale, safeguarded geometry configurator rather than as a claim of universally lower computational cost. The 5% illumination floor is inactive at the nominal 3-dB proposed optimum, where the minimum achieved illumination is 76.6% of the AP budget. For reference, increasing the total Tx covariance dimension from 9 to 18 raises median open-source SDP solve time from 0.366 to 2.144 s.

For the revised calibration-aware solver with $N_{\text{sc}} = 128$ training scenarios, the 2° mixed orientation/actuator experiment checks 26 Tx/Rx endpoint candidates and terminates after one sweep with no accepted move. The 26 robust candidate solves require 357.2 s of aggregate solver time and 398 s wall-clock time, with peak memory below 1 GB. The orientation-only counterpart requires 622.6 s aggregate solver time and 695 s wall-clock time, while the actuator-only counterpart requires 352.4 s aggregate solver time for the same 26-candidate sweep. Thus the tested robust sweeps retain the same outer candidate count but span 352.4–622.6 s of aggregate inner-solver time as the nuisance structure changes. Together with the nominal 3-dB profile above (39 exact fixed-shape SDP solves in 7.54 s), these measurements support the claimed linear dependence of the outer work on the number of tested endpoints while showing that the scenario-expanded fixed-shape SDP dominates practical runtime. They are implementation-specific measurements, not a universal complexity law.

8.3. Minimum-spacing closeout and stress test

Equation (6) already gives an analytic closeout for the nominal hardware setting: because the fixed coordinate separation is $\lambda/2$ and the morphing coordinate is orthogonal, every feasible shape satisfies $d_{\min} \leq \lambda/2$ automatically. We nevertheless stress-test stricter hardware requirements at $\Gamma = 3$ dB using $d_{\min}/\lambda \in \{0.50, 0.60, 0.75\}$. The fixed-shape covariance SDP is unchanged; the only added mechanism is the explicit geometry-feasibility gate in Equation (38h), which rejects an infeasible endpoint before any SDP solve.

Figure 6a reports the resulting position CRBs. At $d_{\min} = 0.5\lambda$, the canonical zero initialization reproduces the nominal CRB of 1.2893 m^2 , with a minimum pair distance of exactly 0.5λ . Since this initialization becomes infeasible under stricter spacing, two deterministic mirrored feasible starts are considered for each higher threshold. All four runs satisfy the spacing, SINR, and power constraints. Their attained CRBs range from 2.3350 to 2.4877 m^2 , indicating higher localization error bounds under the tested strict-spacing configurations. Since the initial geometries also differ, these results are interpreted as a feasibility stress test rather than an isolated measure of the spacing-induced CRB penalty.

Figure 6b further explains the feasibility mechanism. For both $d_{\min} = 0.6\lambda$ and 0.75λ , the final minimum pair distance reaches 1.0λ . Each strict-spacing run rejects 13 infeasible endpoint candidates and accepts no geometry updates. The identical outcomes arise because both thresholds retain the same feasible endpoint neighborhood under the finite candidate search. This does not imply equivalence of the continuous feasible sets or establish a universal monotonic CRB dependence on $d_{\min}$. Instead, the experiment verifies that the proposed search enforces stricter hardware spacing without violating the communication or power constraints.




Figure 6. Minimum-spacing closeout and strict hardware stress test at $\Gamma = 3$ dB. **(a)** The canonical 0.5λ setting is automatically feasible and reproduces the nominal result, while stricter spacing requirements reduce the available geometry freedom and increase the attained CRB for the registered feasible starts; **(b)** Under 0.6λ and 0.75λ , infeasible endpoint candidates are rejected before the covariance SDP; both tested thresholds leave the same surviving finite endpoint neighborhood, so their coincident outcomes are not interpreted as equivalence of the continuous constraints.

8.4. Illumination-floor sensitivity

To verify that the default illumination floor is a regularity safeguard rather than a performance-tuned sensing target, we repeat the proposed design for $q_t/P_t \in \{0.01, 0.02, 0.05, 0.10\}$ under both the nominal $\Gamma = 3$ dB load and a communication-stressed $\Gamma = 9$ dB load. Only the illumination-floor fraction is changed; all geometry, covariance, SINR, power, and endpoint-search settings are retained. As a provenance check, the 5% points exactly reproduce the corresponding Stage-14 common-SINR CRBs at both 3 and 9 dB.

At 3 dB, the floor is inactive throughout the full 1%–10% sweep: the minimum optimized illumination remains 76.60% of the AP power budget, while the CRB stays within 1.28932713–1.28932722 m². The maximum change relative to the 5% setting is only $3.66 \times 10^{-6}\%$, confirming that the nominal result is not created by the chosen 5% floor. At 9 dB, by contrast, all four floors become active and the minimum achieved illumination tracks the imposed 1, 2, 5, and 10% levels. The corresponding CRBs are 1.6781, 1.6787, 1.6995, and 1.7029 m², respectively. Relative to the default 5% case, the changes are -1.26% , -1.22% , and $+0.20\%$ for 1, 2, and 10%, and the total CRB spread over the entire 1%–10% range is 1.48%. Thus the floor is nonbinding at the nominal operating point and only mildly performance-relevant under the high communication load, where it exposes the expected communication–sensing resource tradeoff while preserving nuisance-block regularity. Figure 7 shows that the illumination floor is effectively inactive at the nominal load but becomes active under the stressed 9-dB load.




Figure 7. Sensitivity to the illumination regularity floor. **(a)** At the nominal 3-dB communication load, the optimized CRB is effectively unchanged over the 1%–10% floor sweep, whereas at 9 dB the floor becomes active and produces only a mild CRB variation; **(b)** The 3-dB design naturally maintains about 76.6% minimum illumination, while the 9-dB solutions track the imposed floor, confirming that the default 5% value is nonbinding at the nominal operating point.

8.5. Link-level communication validation

To validate communication reliability beyond the optimization-level SINR constraints, we perform an uncoded quadrature phase-shift keying (QPSK) link-level replay using the frozen optimized designs from the common-SINR sweep. No geometry or covariance is redesigned for this test. For each target $\Gamma \in \{-3, 0, 3, 6, 9\}$ dB, the recovered rank-one communication beams are held fixed, the optimized probing covariance is explicitly realized as a Gaussian sensing waveform and therefore enters the received signal as actual interference, and 10^6 QPSK symbols are simulated per user. Multiuser interference and receiver noise are included in the same realization. Before Monte Carlo evaluation, each replay is provenance-locked to the corresponding Stage-14 design through the nuisance-aware CRB and post-recovery SINR checks. All five operating points pass these consistency gates; the largest relative SINR-target shortfall is 1.82×10^{-6} at 9 dB.

Figure 8 reports the pooled bit error rate (BER) across the two users with Wilson 95% confidence intervals. The BER decreases monotonically from 2.391×10^{-1} at -3 dB to 1.591×10^{-1} , 7.885×10^{-2} , 2.287×10^{-2} , and 2.383×10^{-3} at 0, 3, 6, and 9 dB, respectively. The two users remain closely matched at every operating point; for example, their 9-dB BERs are 2.371×10^{-3} and 2.395×10^{-3} . The Monte Carlo curve closely follows the coherent-QPSK Gaussian-equivalent benchmark, confirming that the SINR-constrained joint ISAC designs translate consistently into link-level communication reliability even when the optimized sensing waveform is physically realized as communication interference. This experiment is a post-design validation of the existing SINR-constrained formulation rather than a BER-oriented redesign.




Figure 8. Unencoded QPSK link-level replay of the frozen optimized ISAC designs. The probing covariance is explicitly realized as communication interference, together with multiuser interference and receiver noise. With 10^6 QPSK symbols per user and operating point, the pooled Monte-Carlo BER closely follows the Gaussian-equivalent coherent-QPSK benchmark over the full -3 to 9 dB SINR range.

8.6. Physical 28-GHz reference and gain attribution

The physical evaluation retains the three dual-functional AP locations, $L = 64$, and the same displacement envelope. Communication links follow Friis scaling with a 20-dB excess-loss margin; with target radar cross section (RCS) σ_{RCS} , the mean bistatic echo gain is

$$\beta_{tr}^s = \frac{G_t G_r \lambda^2 \sigma_{\text{RCS}}}{(4\pi)^3 r_t^2 r_r^2 L_{\text{sys}}}. \quad (48)$$

The receiver noise is 4.013×10^{-13} W (-93.96 dBm). Power gains are normalized internally by this noise level for numerical conditioning without changing the SINR or EFIM. Table 4 lists the physical link-budget parameters and the perturbations used in the paired tests.

Six methods are evaluated on the same 48 channel/geometry realizations. Shapes are designed at the nominal geometry and frozen before the test bank is opened; communication and sensing covariances are then reoptimized per realization. All methods are feasible in all 48 trials. Table 5 reports the resulting physical-scale position-CRB statistics with 95% confidence intervals (CIs).

Table 4. Physical link budget and paired-test perturbations.

Parameter	Value
Carrier/bandwidth	28 GHz/20 MHz
Per-AP power/snapshots	30 dBm/64
Tx/Rx element gain	5 dBi/5 dBi
Noise figure/noise power	7 dB/−93.96 dBm
Communication excess loss	20 dB
Radar system loss/mean RCS	6 dB/1 m ²
Communication/radar Rician K	8 dB/6 dB
Communication/radar shadowing standard deviation	3 dB/2 dB
User jitter/target region	5-m standard deviation/12-m disk

Table 5. Physical position CRB over 48 paired trials.

Method	Mean (m ²)	95% CI (m ²)
Fixed + isotropic	64.18	[57.85, 70.79]
Covariance only	27.57	[25.09, 30.22]
Random vertex	12.00	[10.81, 13.25]
Rx-only morphing	9.13	[8.24, 10.06]
Tx-only morphing	12.07	[10.94, 13.28]
Joint morphing	7.77	[7.04, 8.56]

Figure 9 visualizes the same six-method comparison with 95% bootstrap intervals.




Figure 9. Mean physical-scale position CRB with 95% bootstrap intervals over 48 paired trials. Joint morphing attains the lowest mean CRB, while the same paired trial bank is used for all six designs.

Joint morphing beats covariance-only in all 48 paired trials. For trial i we use the paired log-gain

$$G_i = 10 \log_{10} \left(\frac{\text{CRB}_{\text{cov},i}}{\text{CRB}_{\text{joint},i}} \right). \quad (49)$$

Averaging G_i gives 5.512 dB with a 95% bootstrap interval of [5.375, 5.652] dB. This quantity is deliberately interpreted as a *combined* geometry-activation and optimization gain: moving from the collapsed fixed geometry to the full joint design both occupies the permitted aperture envelope and selects a favorable shape within it. The archived single random full-envelope vertex provides only a first same-envelope control, for which the paired gain is 1.86 dB. The ensemble experiment below is the stronger attribution test.

To test out-of-sample persistence without selecting shapes by performance, the first 20 frozen random-shape indices are fixed in advance and evaluated on the same 48 physical realizations, yielding 960 paired shape–trial comparisons. Joint morphing wins all 960. Across the 20 shape means, the mean and median gains are 2.78 and 2.65 dB, and the 5th–95th percentile range is [1.74, 4.42] dB. Figure 10 shows the corresponding per-shape mean out-of-sample (OOS) gains and bootstrap intervals.




Figure 10. Mean paired joint gain for 20 prespecified random full-envelope geometries over the shared 48-trial bank. Error bars denote 95% bootstrap intervals of each shape mean.

8.7. Geometry generalization: random apertures and spatial redesign

We first replace the single random control by 100 independently frozen Tx/Rx endpoint geometries. After each shape is fixed, communication and probing covariances are reoptimized under the same power, SINR, and illumination constraints. The joint reference CRB is 7.0456 m²; the random ensemble has mean and median CRBs of 15.13 and 14.28 m². Joint morphing beats all 100 samples, with mean/median gains of 3.15/3.07 dB and a 5th–95th percentile range of [1.55, 5.28] dB. Even the best random geometry remains 0.77 dB worse. Figure 11 shows the empirical cumulative distribution function (CDF) of the gain over the 100 independently frozen random geometries.




Figure 11. Joint gain relative to 100 independently frozen full-envelope random geometries, with covariance reoptimized after each shape is fixed.

Spatial generalization is evaluated in two complementary ways. First, the source-locked nominal covariance-only and joint policies are transferred without redesign over a dense 61×61 grid. Joint remains better at 93.95% of the 3721 locations, with mean/median gains of 5.07/5.27 dB; the small negative tail therefore quantifies policy mismatch rather than being silently discarded. Second, both methods are freshly redesigned at every point of an 11×11 grid over $[-30, 30]^2$ m. Covariance-only and one-sweep joint optimization start from the same zero-displacement geometry at each location. All 121 solutions are finite and joint morphing wins at all 121 points, with mean/median gains of 5.79/5.68 dB, a 5th–95th percentile range of [3.40, 8.32] dB, and an observed range of [2.37, 10.80] dB. The median per-point runtime, including the repeated covariance resolves in the one-sweep shape search, is 10.72 s on Ibex. Figures 12 and 13 show, respectively, the location-wise gain map and its empirical distribution over all 121 reoptimized target locations.




Figure 12. Reoptimized spatial CRB gain. Covariance-only and one-sweep joint morphing are independently redesigned at each of the 121 target locations; positive gain is obtained throughout the grid.




Figure 13. Empirical CDF of the 121 reoptimized spatial gains. The median is 5.68 dB and the 5th–95th percentile interval is [3.40, 8.32] dB.

The random-shape and spatial experiments close two different attribution gaps. The 100/960 comparisons show that the optimized aperture is not favored merely because one unlucky random geometry was chosen, while the 121-point redesign shows that the advantage is not confined to the nominal target location.

8.8. Source-locked finite-sample ML validation

We next test whether the CRB ranking survives finite samples. The formal variable-projection ML protocol fixes SNR offsets $\{-5, 0, 5, 10, 15\}$ dB, 128 replicates per method/SNR point, common random channel/noise draws across methods, a ± 20 -m search box, a 4-m coarse grid, and five limited-memory Broyden–Fletcher–Goldfarb–Shanno with bounds (L-BFGS-B) starts. Outage is defined as Euclidean position error above 10 m, and failed or boundary estimates are retained.

The four designs are fixed before ML evaluation and checked against the archived design results. All 2560 estimation cases complete without optimizer failure. Figures 14 and 15 show the localization RMSE relative to the position error bound (PEB) and the estimator outage, respectively. Table 6 summarizes ML efficiency at high SNR.

At 5/10/15 dB, the joint RMSE–RMS-PEB gap has magnitude below 0.45 dB and joint RMSE is lower than covariance-only at every checkpoint. At -5 dB, joint morphing records 6.654-m RMSE and 10.94% outage, compared with 15.027 m and 59.38% for covariance-only, making the threshold-region separation explicit. The slightly negative high-SNR empirical gaps are not interpreted as beating the CRB: the plotted RMSE is a finite Monte-Carlo realization of a mildly biased nonlinear estimator, and post-run bootstrap intervals of the mean-square-error (MSE)/mean-CRB gap include zero at all three checkpoints.




Figure 14. Source-locked variable-projection ML validation. Joint morphing attains the lowest estimator RMSE over the full frozen SNR sweep; at 5–15 dB, its RMSE remains within 0.45 dB of the corresponding joint RMS-PEB.




Figure 15. Estimator outage, defined as position error above 10 m. Joint morphing has the lowest observed outage at −5 dB and no observed outage over 128 trials per point from 0 to 15 dB; the low-SNR threshold region is retained rather than removed by post-hoc SNR selection.

Table 6. Exact ML efficiency at the frozen high-SNR checkpoints.

SNR (dB)	Joint RMSE (m)	Joint RMS-PEB (m)	Gap (dB)	Covariance-Only RMSE (m)
5	1.481	1.532	−0.293	4.023
10	0.822	0.861	−0.409	1.899
15	0.460	0.484	−0.444	1.029

8.9. Calibration boundary and regime transition

The positive nominal results above assume known post-calibration geometry. The isolated 2° negative control reported in Section 5. shows a 3.23-dB tail penalty and essentially no recovery from covariance reallocation. To test whether that conclusion is specific to one uncertainty level, we next use a disjoint common-random-number sweep. In this sweep the standardized orientation draws are rescaled while the archived element-error draws are held fixed; consequently, its 0° point is the zero-orientation baseline for the sweep and is not numerically identical to the separate known-geometry reference used for the 3.23-dB statement. All reported sweep values are risk summaries of local conditional CRBs, not the exact hybrid CRB in Equation (32). Throughout this subsection, the hybrid-SAA CRB is reported only for the calibration-aware design itself, whereas the displayed mean/95th-percentile/CVaR₉₅ numbers are computed after freezing the design and evaluating scenario-wise local CRBs. Thus the two columns are complementary design and deployment-risk diagnostics rather than interchangeable CRB estimators.

The orientation component is swept over

$$\sigma_\delta \in \{0, 0.25^\circ, 0.5^\circ, 1^\circ, 2^\circ, 3^\circ, 5^\circ\}. \quad (50)$$

The 2° point reproduces the archived nominal and robust values within 2.4×10^{-5} relative error. Across the sweep, nominal CVaR₉₅ rises from 8.08 m^2 at 0° to 10.72, 18.35, 31.03, and 71.67 m^2 at $1^\circ, 2^\circ, 3^\circ$, and 5° , respectively. Yet calibration-aware covariance reallocation changes CVaR₉₅ by only -0.051 to $+0.021$ dB relative to the nominal covariance under the same nuisance model.

We additionally place calibration uncertainty directly inside the joint geometry/covariance optimization of Section 7.4. The design uses 128 common-random-number training scenarios and is evaluated only after freezing the design on an independent 4096-scenario bank. Three mechanism-separated cases are tested. First, with orientation uncertainty only at $\sigma_\delta = 2^\circ$, all 26 admissible Tx/Rx endpoint candidates are examined and none is accepted. The robust-covariance and robust-joint designs are therefore identical, with an out-of-sample hybrid-SAA CRB of 16.8074 m^2 and CVaR₉₅ local CRB of 17.4272 m^2 . Second, with actuator execution error only at $\sigma_e = 0.05\lambda$, the search again accepts no endpoint move; both robust designs give a hybrid-SAA CRB of 7.3259 m^2 and CVaR₉₅ of 8.0371 m^2 . Third, with $\sigma_\delta = 2^\circ$ and $\sigma_e = 0.05\lambda$ jointly, no endpoint move is accepted and both robust designs give a hybrid-SAA CRB of 17.0937 m^2 and CVaR₉₅ of 18.2017 m^2 . Hence the nominal commanded aperture is coordinatewise endpoint-stable in all three tested calibration regimes; the large mixed-error degradation is dominated by orientation uncertainty, whereas the actuator-only hybrid CRB remains close to the 7.0456 m^2 nominal value.

As an additional claim-to-proof check, we freeze the D1/D2 geometries and reoptimize their covariances under a stricter 29-dimensional nuisance model that retains target position, the real/imaginary parts of the bistatic gains, three orientation biases, and Tx/Rx differential actuator modes. The frozen D1 and D2 geometries are identical to machine precision, and the strict audit returns exactly the same mean-CVaR objective (17.5083 m^2) and CVaR₉₅ (17.8192 m^2) for both, with a maximum LMI/direct-CRB gap of 1.01×10^{-5} . This frozen-design audit is a consistency check of the full nuisance treatment; it is not presented as a separate 29-dimensional outer geometry search. Together, these results show that calibration error is now included *inside* the geometry optimization rather than assessed only a posteriori, while also

confirming that endpoint shape reconfiguration cannot recover the dominant orientation-information loss predicted by Equation (30) in the tested regime. Figure 16 visualizes the rapid growth of orientation-induced tail risk and the near-overlap of the nominal and calibration-aware covariance designs.

Figure 17 summarizes the corresponding design interpretation by combining the angular-diversity factor with the scalar information ceiling in Equation (30). Small angular separation is geometry-limited; moderate separation and mild orientation uncertainty are design-controllable; and large residual orientation uncertainty produces a calibration-limited regime in which additional covariance tuning has little leverage.

Taken together, the 0° – 5° sweep and the matrix hybrid EFIM support a sharper boundary statement than the original 2° negative control: geometry reconfiguration is effective after calibration, but probing covariance cannot reconstruct a geometric reference that has itself become uncertain.




Figure 16. Orientation-uncertainty sweep with common random numbers. The tail risk grows rapidly with σ_δ , whereas nominal and calibration-aware covariance designs remain nearly indistinguishable.




Figure 17. Analytical geometry-calibration regime map. Background colors denote geometry-limited, design-controllable, and calibration-limited regions; contours show a normalized effective-information score; s.d. denotes standard deviation.

9. Design implications

The derived bounds and frozen diagnostics support five system-level design implications.

- (1) Waveform rank and aperture are complementary resources. A reconfigurable aperture does not by itself guarantee AoD information: when the effective probing covariance is rank one, the transmit derivative remains in the nuisance-removed signal span. Shape control is therefore most valuable when paired with nondegenerate probing structure.
- (2) Strong local information does not guarantee strong localization. The position EFIM depends on both local information intensity and the cross-direction determinant term. Even aggressive local aperture shaping cannot compensate for nearly parallel network bearings.
- (3) Geometry control cannot recover a lost geometric reference. Once residual orientation uncertainty dominates the hybrid information budget, the effective information saturates and covariance-only robustification offers very limited recovery. The full 0° – 5° sweep shows that this is a regime-level phenomenon rather than an isolated 2° operating point.
- (4) The optimized shape advantage is not tied to one random geometry or one target point. The 100-shape ensemble and 20-shape physical-bank test show that the joint solution consistently dominates randomly occupied full-envelope geometries. The 121-point reoptimized map further shows a positive local redesign advantage throughout the tested target region.
- (5) The CRB ranking is estimator-relevant above threshold. The source-locked ML replay preserves the joint-versus-covariance ordering at all preregistered high-SNR points while explicitly exposing the low-SNR threshold region. The bound is therefore used as a design predictor, not as a substitute for finite-sample validation.

Together, these implications distinguish three different failure modes that should not be conflated: insufficient probing rank, poor network angular conditioning, and loss of geometric reference through calibration uncertainty. The appropriate remedy is different in each regime.

10. Bound interpretation and limitations

All nominal sensing values above are position CRBs in square metres, not estimator RMSEs. For a locally unbiased estimator under the adopted regular model,

$$\begin{aligned} \mathbb{E}[\|\hat{\mathbf{p}} - \mathbf{p}\|^2] &\geq \text{tr}(\mathbf{I}_{\mathbf{p}}^{-1}) \triangleq \text{CRB}(\mathbf{p}), \\ \text{PEB} &= \sqrt{\text{CRB}(\mathbf{p})}. \end{aligned} \tag{51}$$

The CRB is reported in m^2 , whereas PEB and RMSE are reported in m. Section 8.8. supplies one explicit finite-sample estimator check: the estimator, target, SNR grid, search region, Monte-Carlo bank, failure rule, and outage definition are all frozen before the exact replay. Its conclusions should therefore be read within that operating regime rather than as a global claim of CRB attainability.

The nominal optimization is local to a coarse target location and assumes known post-calibration geometry. The revised calibration-aware branch also optimizes commanded geometry and probing covariance over sampled orientation and actuator uncertainty, but its endpoint-stability certificate remains

local to the adopted discrete shape-search set. Orientation and actuator uncertainty are separately treated in the augmented information model. The independent bistatic gains form a conservative angle-only model; a trusted shared reflectivity or propagation law can introduce additional range and cross-link information and requires a new joint EFIM.

The signal model also assumes coherent synchronization, resolvable single-bounce returns, quasi-static shapes within a sensing frame, and a far-field two-dimensional angle model. Near-field propagation, mutual coupling, unresolved multi-target returns, clock offsets, and dynamic tracking [30] must enter the complete parameter vector before nuisance elimination and cannot be inferred by simply rescaling the present bound.

11. Conclusion

This paper developed a nuisance-aware framework for geometry-reconfigurable cell-free ISAC. Eliminating independent complex bistatic gains isolates the steering variations that remain identifiable, yields the exact rank-one probing limitation, and separates local aperture sensitivity from network angular diversity. The comparison with a known-gain/raw-derivative surrogate shows that nuisance omission can create unbounded CRB optimism by turning a non-identifiable Tx-angle direction into apparently positive information. We also make the architectural scope explicit: the EFIM itself extends to distributed multistatic sensing, whereas the joint optimum is specific Cell-Free-ISAC through user-centric serving masks, per-AP power budgets, multiuser SINR constraints, and probing/communication resource competition. A signal-consistent joint covariance/shape design then exploits these quantities: the fixed-shape block is convex on the finite-CRB domain, communication rank relaxation is exactly recoverable under dedicated probing, and the revised strict safeguarded updates terminate finitely at a coordinatewise endpoint-stable geometry. A calibration-aware hybrid-SAA branch further places orientation uncertainty and sampled actuator execution errors inside the joint geometry/covariance design while preserving an SDP fixed-shape block.

The physical and journal-scale evidence is mutually consistent. Under a common-system algorithm-level comparison, the safeguarded endpoint search has the lowest nuisance-aware CRB at all five tested SINR targets and all three aperture envelopes, and wins all 20 paired random-initialization trials against Qin-style AO-SCA, Chen-style MM-RGP, and Lou-style Taylor-AO updates. A frozen-design QPSK replay further verifies the communication side of this tradeoff: with the optimized probing waveform realized as actual interference, the pooled uncoded BER decreases from 2.391×10^{-1} at -3 dB to 2.383×10^{-3} at 9 dB and closely follows the Gaussian-equivalent coherent-QPSK benchmark. A minimum-spacing audit further confirms that $d_{\min} \leq 0.5\lambda$ is guaranteed by the adopted orthogonal morphing geometry, while stricter 0.6λ and 0.75λ stress cases are enforced by explicit pre-SDP candidate rejection with zero spacing violation, at the expected cost of reduced geometry freedom. An illumination-floor sweep further shows that the default 5% value is inactive at the nominal 3-dB operating point, where the optimized minimum illumination is 76.6%; under the 9-dB stress load the floor becomes active, but the CRB varies by only 1.48% across the full 1%–10% sweep. Computationally, the safeguarded search has a finite and explicit outer solve count at most $2N_{\text{Tx}}$ additional fixed-shape SDPs per nominal sweep and $2(N_{\text{Tx}} + N_{\text{Rx}})$ robust SDPs per calibration-aware sweep, while measured runtime confirms that the scenario-expanded inner SDP, rather than endpoint bookkeeping, is the dominant cost. The convergence claim is correspondingly

limited to finite termination and coordinatewise endpoint stability on the registered endpoint set; no continuous-domain KKT or global-optimality claim is made. Likewise, the hybrid CRB and scenario-wise conditional-CRB CVaR are kept separate: the former averages complete information before nuisance elimination, whereas the latter is a frozen-design deployment-risk summary. The 48-trial 28-GHz campaign gives a 5.51-dB combined joint-over-covariance-only CRB gain; 100 random full-envelope geometries and 960 paired random-shape/out-of-sample comparisons show that the optimized aperture advantage is not tied to a single random control; and fresh redesign retains positive gain at all 121 spatial test points. The source-locked ML replay further shows that the CRB ranking predicts finite-sample performance above threshold, with zero optimizer failures over 2560 cases and a joint RMSE–RMS-PEB gap below 0.45 dB at the 5/10/15-dB checkpoints. The calibration sweep provides the complementary limit: orientation uncertainty drives CVaR₉₅ local CRB from 8.08 to 71.67 m² over 0°–5°, while covariance redistribution changes that tail metric by at most 0.051 dB. Geometry reconfiguration is therefore a strong post-calibration sensing resource, but it cannot substitute for the geometric reference that calibration provides. Consistently, the revised calibration-aware joint search reaches an endpoint-stable commanded geometry with no accepted Tx/Rx endpoint move in the orientation-only ($\sigma_\delta = 2^\circ$), actuator-only ($\sigma_e = 0.05\lambda$), and mixed-error audits. A frozen-design 29-dimensional nuisance check further gives identical D1/D2 mean–CVaR objectives and CVaR₉₅ values, confirming that the conclusion is not an artifact of a purely posteriori calibration assessment.

Appendix A. Complete real-FIM elimination of one complex gain

For one link, write $\mathbf{y} = \alpha s(\vartheta) + \mathbf{n}$ with $\mathbf{n} \sim \mathcal{CN}(\mathbf{0}, \mathbf{C})$ and $\alpha = \alpha_R + j\alpha_I$. Let $\mathbf{D} = \partial s / \partial \vartheta^T$. For a complex Gaussian model with covariance independent of the real parameter vector, $[\mathbf{J}]_{ij} = 2 \operatorname{Re}\{(\partial_i \mu)^H \mathbf{C}^{-1} (\partial_j \mu)\}$. The complete derivatives are $\partial_\vartheta \mu = \alpha \mathbf{D}$, $\partial_{\alpha_R} \mu = s$, and $\partial_{\alpha_I} \mu = js$. Eliminating (α_R, α_I) gives

$$\mathbf{J}_e(\vartheta) = 2|\alpha|^2 \operatorname{Re} \left\{ \mathbf{D}^H \mathbf{C}^{-1/2} \Pi_{\mathbf{C}^{-1/2} \mathbf{s}}^\perp \mathbf{C}^{-1/2} \mathbf{D} \right\}, \quad (52)$$

which reduces to Equation (12) for white noise. This proof fixes the factor of two and makes explicit that the complex gain is represented by two real nuisance coordinates.

Appendix B. Finite-CRB geometry identity

Let $\mathbf{u}_m = [-\sin \theta_m, \cos \theta_m]^T$ and $\mathbf{I}_p = \sum_m w_m \mathbf{u}_m \mathbf{u}_m^T$. Writing the two rows of the $2 \times M$ weighted direction matrix as $\mathbf{r}_1$ and $\mathbf{r}_2$, Cauchy–Binet gives

$$\det \mathbf{I}_p = \sum_{m < m'} w_m w_{m'} \det[\mathbf{u}_m, \mathbf{u}_{m'}]^2, \quad (53)$$

and $\det[\mathbf{u}_m, \mathbf{u}_{m'}] = \sin(\theta_{m'} - \theta_m)$, yielding Equation (36). Therefore finite two-dimensional localization requires at least two nonparallel informative directions.

Appendix C. Why hybrid and conditional risk bounds differ

Let γ be a random calibration variable. The hybrid Fisher information is an expectation of the complete information matrix plus prior information. Schur complementation is nonlinear, so in general

$$\text{Schur}(\mathbb{E}_\gamma[\mathbf{J}(\gamma)]) \neq \mathbb{E}_\gamma[\text{Schur}(\mathbf{J}(\gamma))]. \quad (54)$$

Moreover, neither expression is the arithmetic mean of $\text{tr}(\mathbf{I}_p(\gamma)^{-1})$. The latter remains useful as a scenario-wise deployment-risk statistic (e.g., the 95th percentile or CVaR), but it should not be labeled the exact hybrid CRB.

Appendix D. Diagnostic mean–CVaR calibration program

For completeness, the calibration-aware covariance branch used in the frozen negative control is formulated as a diagnostic risk program, not as evidence of a successful robustness gain. Using the same N_{sc} scenario-count notation as in the calibration-aware design, for calibration scenario s , let $\mathbf{H}_s$ be the complete prior-regularized real FIM for position and all nuisance variables, and let $\mathbf{E}_p = [\mathbf{I}_2, \mathbf{0}]$ select position. For fixed commanded shapes, each entry of $\mathbf{H}_s$ is affine in the probing covariances. Introducing $\mathbf{Z}_s$, c_s , ζ , and u_s gives

$$\begin{aligned} \begin{bmatrix} \mathbf{Z}_s & \mathbf{E}_p \\ \mathbf{E}_p^T & \mathbf{H}_s \end{bmatrix} \succeq \mathbf{0}, \quad c_s \geq \text{tr}(\mathbf{Z}_s), \\ u_s \geq c_s - \zeta, \quad u_s \geq 0. \end{aligned} \quad (55)$$

At CVaR level ρ , the objective

$$(1 - \omega) \frac{1}{N_{\text{sc}}} \sum_s c_s + \omega \left(\zeta + \frac{1}{(1 - \rho)N_{\text{sc}}} \sum_s u_s \right) \quad (56)$$

remains an SDP together with the original communication, power, and illumination constraints. Convexity of this diagnostic program does not imply that the available covariance control authority is sufficient to overcome a calibration ceiling; the frozen results in Section 8. show that it is not sufficient in the tested orientation regime.

Appendix E. Element-error and differential-mode boundary

Equation (33) is obtained by forming the local joint information for the bearing and differential element errors, adding the independent Gaussian actuator prior, and eliminating the actuator block. The projector $\mathbf{P}_m^\perp$ removes the common displacement direction exactly. When the actuator covariance is correlated or multiple mechanical modes are coupled, the scalar expression must be replaced by the complete matrix Schur complement.

A higher-dimensional differential-mode calibration model can be handled by the same complete-FIM construction, but its usefulness depends on a physically justified mechanical covariance model. We therefore keep the closed-form independent-error ceiling above as the reproducible analytical boundary

and do not promote an unfinished differential-mode optimization study as manuscript evidence.

Data availability statement

The numerical evidence required to reproduce the reported aggregate statistics is provided with the supplementary submission package. Additional implementation files are available from the corresponding author upon reasonable request.

Declaration of generative AI and AI-assisted technologies

During manuscript preparation, the author used ChatGPT (OpenAI) for language polishing, readability improvement, and LaTeX formatting assistance. Numerical results were produced by simulation and optimization scripts; the scientific claims, equations, statistics, and final manuscript content were reviewed and verified by the author. No generative image model was used to create the scientific plots in this manuscript. The author takes full responsibility for the integrity, originality, and accuracy of the work.

Acknowledgments

The author acknowledges the research and computing environment provided by King Abdullah University of Science and Technology (KAUST).

Conflicts of interest

The author declares no conflicts of interest.

References

- [1] Demirhan U, Alkhateeb A. Cell-Free ISAC MIMO systems: joint sensing and communication beamforming. *IEEE Trans. Commun.* 2025, 73(6):4454–4468.
- [2] Behdad Z, Demir OT, Sung KW, Björnson E, Cavdar C. Multi-static target detection and power allocation for integrated sensing and communication in cell-free massive MIMO. *IEEE Trans. Wirel. Commun.* 2024, 23(9):11580–11596.
- [3] Galappaththige D, Tellambura C. Cell-free integrated sensing and communication. *Found. Trends Netw.* 2025, 16(1–2):1–170.
- [4] Wong K, Shojaeifard A, Tong K, Zhang Y. Performance limits of fluid antenna systems. *IEEE Commun. Lett.* 2020, 24(11):2469–2472.
- [5] Wong K, Shojaeifard A, Tong K, Zhang Y. Fluid antenna systems. *IEEE Trans. Wirel. Commun.* 2021, 20(3):1950–1962.
- [6] Chai Z, Wong K, Tong K, Chen Y, Zhang Y. Port selection for fluid antenna systems. *IEEE Commun. Lett.* 2022, 26(5):1180–1184.

- [7] Zhu L, Ma W, Zhang R. Movable antennas for wireless communication: opportunities and challenges. *IEEE Commun. Mag.* 2024, 62(6):114–120.
- [8] Zhu L, Ma W, Mei W, Zeng Y, Wu Q, *et al.* A tutorial on movable antennas for wireless networks. *IEEE Commun. Surv. Tutor.* 2026, 28:3002–3054.
- [9] Zhu L, Ma W, Zhang R. Modeling and performance analysis for movable antenna enabled wireless communications. *IEEE Trans. Wirel. Commun.* 2024, 23(6):6234–6250.
- [10] Lyu W, Yang S, Xiu Y, Zhang Z, Assi C, *et al.* Movable antenna enabled integrated sensing and communication. *IEEE Trans. Wirel. Commun.* 2025, 24(4):2862–2875.
- [11] Chen L, Zhao M, Zhao M, Zhang R. Antenna position and beamforming optimization for movable antenna enabled ISAC: optimal solutions and efficient algorithms. *IEEE Trans. Signal Process.* 2025, 73:3812–3828.
- [12] Ma W, Zhu L, Zhang R. Movable antenna enhanced integrated sensing and communication via antenna position optimization. *IEEE Trans. Signal Process.* 2026, 74:1522–1537.
- [13] Lou X, Xia W, Zhu Y, Wong KK, Chae CB. Multi-target beamforming optimization for fluid antenna-enabled multi-static ISAC. *IEEE Trans. Cogn. Commun. Netw.* 2026, 12:1382–1394.
- [14] Xiu Y, Zhao Y, Yang R, Lyu W, Niyato D, *et al.* Robust optimization for movable antenna-aided cell-free ISAC with time synchronization errors. *IEEE Trans. Wirel. Commun.* 2026, 25:10082–10097.
- [15] Xiu Y, Zhao Y, Yang R, Dong Z, Lyu W, *et al.* Distributed distortion-aware robust optimization for movable antenna-aided cell-free ISAC systems. *arXiv* 2025, arXiv:2508.13839.
- [16] Zheng Q, Zhu P. Multi-static ISAC assisted by double-side fluid antenna system. *arXiv* 2026, arXiv:2604.25234.
- [17] He C, He J, Abdallah A, Eltawil AM. Localization bounds and shape design for cell-free ISAC with fluid intelligent metasurfaces. *arXiv* 2026, arXiv:2608.16634.
- [18] Kay SM. *Fundamentals of Statistical Signal Processing: Estimation Theory*, 1st ed. Englewood Cliffs: Prentice Hall, 1993.
- [19] Trees HLV. *Detection, Estimation, and Modulation Theory*, 1st ed. New York: Wiley, 2001.
- [20] Noam Y, Messer H. Notes on the tightness of the hybrid Cramér–Rao lower bound. *IEEE Trans. Signal Process.* 2009, 57(6):2074–2084.
- [21] Haimovich AM, Blum RS, Cimini LJ. MIMO radar with widely separated antennas. *IEEE Signal Process. Mag.* 2008, 25(1):116–129.
- [22] Godrich H, Haimovich AM, Blum RS. Target localization accuracy gain in MIMO radar-based systems. *IEEE Trans. Inf. Theory* 2010, 56(6):2783–2803.
- [23] Shen Y, Win MZ. Fundamental limits of wideband localization—part I: a general framework. *IEEE Trans. Inf. Theory* 2010, 56(10):4956–4980.
- [24] Rockah Y, Schultheiss PM. Array shape calibration using sources in unknown locations—part I: far-field sources. *IEEE Trans. Acoust. Speech Signal Process.* 1987, 35(3):286–299.
- [25] Wan S, Tang J, Zhu W, Zhang N. Identifiability analysis for array shape self-calibration based on hybrid Cramér–Rao bound. *IEEE Signal Process. Lett.* 2014, 21(4):473–477.

- [26] Zhang N, Tang J, Wan S. Identifiability analysis for array shape self-calibration in colocated multiple-input multiple-output radar using Cramér–Rao bound. *IET Signal Process.* 2016, 10(1):81–89.
- [27] Bomfin R, Mezzavilla M, Rangan S, Chafii M. Near field multi-band localization: CRB, efficient estimator, and threshold SNR. *arXiv* 2026, arXiv:2606.12314.
- [28] Boyd S, Vandenberghe L. *Convex Optimization*, 1st ed. Cambridge: Cambridge University Press, 2004.
- [29] Qin H, Chen W, Wu Q, Zhang Z, Li Z, *et al.* Cramér–Rao bound minimization for movable antenna-assisted multiuser integrated sensing and communications. *IEEE Wirel. Commun. Lett.* 2024, 13(12):3404–3408.
- [30] He C, Zhang X, Li G. Twin-in-the-loop optimization and fundamental limits of position–velocity estimation in cell-free ISAC systems. *arXiv* 2026, arXiv:2606.15688.