

# Sub-Nyquist Affine Antenna Switching for Low-Complexity Angle–Doppler Estimation

Alireza Pourafzal\*, Nitin Jonathan Myers<sup>†</sup>, Dario Tagliaferri<sup>‡</sup>, Musa Furkan Keskin\*, and Henk Wymeersch\*

\* Department of Electrical Engineering, Chalmers University of Technology, Sweden

<sup>†</sup> Delft Center for Systems and Control, Delft University of Technology, 2628 CD Delft, Netherlands

<sup>‡</sup> Department of Electronics, Information and Bioengineering, Politecnico di Milano, Milan, Italy

Email: alireza.pourafzal@chalmers.se

**Abstract**—Antenna switching enables low-complexity sensing with a single radio-frequency chain, but target motion couples the switching sequence with Doppler-induced phase evolution. This coupling makes joint two-dimensional direction-of-arrival (DoA) and Doppler recovery rely on a computationally expensive three-dimensional (3D) search. This paper proposes an affine switching strategy for angle–Doppler estimation in integrated sensing and communication arrays. By sequentially switching across affine contours of a planar array, Doppler appears as a deterministic shift of the two-dimensional spatial spectrum in each switching sweep. Multiple sweeps with different slopes then recover the physical DoA and Doppler parameters by linear inversion after per-sweep spatial estimation. Simulations show that, compared with an implemented direct 3D grid-search baseline, the proposed method exhibits a lower signal-to-noise ratio threshold for successful initialization and substantially lowers runtime, while shifted-domain aliasing remains the main limitation in multi-target scenarios.

**Index Terms**—Integrated sensing and communication, affine antenna switching, angle–Doppler estimation, single-RF sensing, beamspace processing.

## I. INTRODUCTION

Large antenna arrays are central to integrated sensing and communication (ISAC), as they provide angular resolution for sensing and spatial selectivity for communication. For low-complexity sensing, a single radio frequency (RF) receiver is desirable, but general analog combining requires phase- or amplitude-controlled networks with nontrivial control, calibration, and insertion-loss overhead [1]. Antenna switching provides a simpler alternative, where the receive chain is connected to one antenna element per sensing snapshot, and the aperture is sampled over time [2]. This switch-based mode is hardware-light, but it requires spatial information to be recovered from sequential receive samples. In this paper, sub-Nyquist refers to acquiring (sampling)  $M$  scalar antenna–time samples from the  $N_x N_y$  receive aperture, rather than observing the complete space–time array data cube.

The main difficulty arises when these switched samples are collected from moving targets. Since different antenna elements are observed at different snapshots, Doppler introduces a phase progression across the switching sequence. If not modeled, this temporal phase is coupled with the direction-of-arrival (DoA)-dependent spatial phase [3], which can degrade conventional DoA estimation in high-mobility settings [4]. Existing angle–Doppler methods typically address

this coupling by adding Doppler as an explicit estimation dimension. Multidimensional spectral methods rely on joint parameter search [5]. Time-division multiplexed multiple-input multiple-output (TDM-MIMO) radars compensate or jointly estimate Doppler-dependent phase shifts in the angle–velocity domain [6]. Repetition-based designs can similarly reduce angle–Doppler coupling by keeping the spatial response fixed over slow-time blocks, but they trade sampling diversity for decoupling rather than optimizing sub-Nyquist array-switching trajectories [7]. Matrix-completion methods recover angle–Doppler structure from incomplete space–time samples [8], while tensor-completion methods extend this idea to higher-dimensional radar data [9]. These approaches are effective, but their search or representation size grows with the angular and Doppler grids. Switching-aware designs show that temporal phase errors can also be shaped by the sampling order. SwiftLink uses contour-based sampling so that carrier-frequency offset (CFO) induces structured beamspace shifts for compressive beam alignment [2], and CFO-robust switching schedules have been studied for DoA estimation [10]. However, these methods do not design integer affine slopes or contour-wise sub-Nyquist samples for two-dimensional DoA and multi-Doppler recovery.

This paper develops an affine antenna-switching framework for low-complexity receive-side angle–Doppler estimation. The main contribution is threefold. First, we show that tying the sampling time affinely to the selected antenna coordinates converts Doppler into slope-dependent shifts of a two-dimensional spatial spectrum. Second, we design feasible integer slopes using Doppler support, wrapping limits, shifted-domain aliasing, and conditioning, together with a spatial prior that improves finite-schedule aperture and coherence. Third, we develop a two-stage estimator that replaces a direct 3D angle–Doppler search with per-sweep 2D searches, cross-sweep association, and linear recovery.

## II. SYSTEM MODEL

We consider an ISAC system with a half-wavelength spaced  $N_x \times N_y$  uniform rectangular array (URA), where  $N = N_x N_y$ . The sensing receiver employs a single RF chain and selects one antenna element per snapshot. Other antennas or RF chains may support communication or additional ISAC functionalities, but they are not considered in this model. We focus

on the narrowband angular–Doppler estimation problem after waveform processing has isolated a delay bin of interest. Delay estimation, range migration, transmit-side processing, clutter, synchronization errors, self-interference, and hardware impairments are outside the present model. This formulation isolates the antenna-switching-induced coupling for joint two-dimensional DoA and Doppler estimation.

#### A. Angle–Doppler Channel and Switched Measurements

Let  $\mathbf{h}_m \in \mathbb{C}^N$  denote the narrowband spatial channel at the full receive array for snapshot  $m \in \{0, \dots, M-1\}$ . The channel contains  $L$  resolvable angular–Doppler paths indexed by  $\ell = 0, \dots, L-1$ . We model the channel as

$$\mathbf{h}_m = \sum_{\ell=0}^{L-1} \alpha_\ell \mathbf{a}_{N_x}(\omega_{x,\ell}) \otimes \mathbf{a}_{N_y}(\omega_{y,\ell}) e^{j\gamma_\ell m}, \quad (1)$$

where  $\alpha_\ell \in \mathbb{C}$  is the  $\ell$ -th path gain,  $(\omega_{x,\ell}, \omega_{y,\ell}) \in [-\pi, \pi)^2$  are the spatial frequencies,  $\gamma_\ell \triangleq 2\pi\nu_\ell T$  is the Doppler-induced phase increment, with Doppler shift  $\nu_\ell$  and the switching interval  $T$ , and  $\mathbf{a}_N(\omega) = [1, e^{j\omega}, \dots, e^{j(N-1)\omega}]^T$ .

The considered ISAC array acquires one spatial measurement per sensing snapshot. Let  $(i_m, j_m) \in \{0, \dots, N_x-1\} \times \{0, \dots, N_y-1\}$  denote the antenna index selected at snapshot  $m$ . The corresponding receive combiner is

$$\mathbf{w}_m = \mathbf{e}_{i_m}^{(N_x)} \otimes \mathbf{e}_{j_m}^{(N_y)} \triangleq \mathbf{e}_{n_m}^{(N)}, \quad n_m = i_m N_y + j_m, \quad (2)$$

with  $\mathbf{e}_i^{(N)}$  as the  $i$ -th canonical basis vector in  $\mathbb{C}^N$ . Thus,  $\mathbf{w}_m$  selects exactly one antenna element. This structure avoids general phase-shifter networks and the associated control overhead of analog combining [2]. The sensing sample is

$$y_m = \mathbf{w}_m^H \mathbf{h}_m + q_m, \quad q_m \sim \mathcal{CN}(0, \sigma^2), \quad (3)$$

with noise variance  $\sigma^2$ . Substituting (1) and (2) into (3) gives

$$y_m = \sum_{\ell=0}^{L-1} \alpha_\ell e^{j(\omega_{x,\ell} i_m + \omega_{y,\ell} j_m + \gamma_\ell m)} + q_m. \quad (4)$$

Hence, each measurement is determined by the selected space–time coordinate  $\mathbf{r}_m \triangleq [i_m, j_m, m]^T$ . The sensing receiver therefore acquires  $M$  samples over the receive array and time, taking one spatial sample per snapshot.

#### B. Problem Statement

For a generic switching set, estimating the parameters in (4) requires joint maximum-likelihood search over  $(\omega_x, \omega_y, \gamma)$  for each path. This 3D search becomes the main computational bottleneck for large arrays, and fine Doppler grids. Our objective is to design the switched sampling set  $\mathcal{R} = \{(i_m, j_m, m)\}_{m=0}^{M-1}$  so that angle–Doppler recovery can be decomposed into lower-dimensional estimation steps while satisfying the single-antenna switching constraint.

### III. AFFINE DOPPLER-AWARE SAMPLING

We now impose structure on the switching set  $\mathcal{R}$  so that the Doppler term in (4) can be absorbed into shifted spatial frequencies. We divide the measurements into blocks, called sweeps. Within each sweep, the selected antenna coordinates and the local snapshot index are constrained by an affine relation. This relation makes the temporal Doppler phase separable into a constant sweep-dependent phase and two




Fig. 1. Affine antenna switching sweeps for a  $6 \times 6$  URA. The gray lines show the affine contours  $p = a_s i + b_s j + c_s$ . The blue markers indicate selected antennas, switched sequentially (one per contour) in increasing order. Doppler manifests as a spatial frequency shift in each sweep. Here  $c_s = 0$ ; in general, nonzero  $c_s$  shifts the feasible contour sequence.

spatial phase terms. Consequently, each sweep admits a two-dimensional spatial spectral representation, while Doppler is encoded through the slope-dependent shifts across sweeps.

#### A. Proposed Affine Sweeps

Let  $\mathcal{I} = \{0, \dots, N_x-1\} \times \{0, \dots, N_y-1\}$  be the antenna-index set. The  $M$  scalar measurement snapshots are partitioned into  $S$  disjoint sweeps  $\{\mathcal{M}_s\}_{s=1}^S$ , where sweep  $s$  contains  $M_s$  snapshots and  $\sum_{s=1}^S M_s = M$ . Let  $p \in \{0, \dots, M_s-1\}$  denote the local snapshot index in sweep  $s$ , and let  $m_s^{(0)} = \sum_{r=1}^{s-1} M_r$  be the first global snapshot index of sweep  $s$ . Then

$$m = m_s^{(0)} + p, \quad m \in \mathcal{M}_s. \quad (5)$$

An affine sweep is specified by an integer triple  $(a_s, b_s, c_s) \in \mathbb{Z}^3$ , with  $(a_s, b_s) \neq (0, 0)$ . For local snapshot index  $p \in \{0, \dots, M_s-1\}$ , the affine rule defines feasible antenna set

$$\mathcal{C}_{s,p} = \{(i, j) \in \mathcal{I} : a_s i + b_s j + c_s = p\}. \quad (6)$$

The sweep is feasible if  $\mathcal{C}_{s,p} \neq \emptyset$  for all  $p = 0, \dots, M_s-1$ . The switching schedule then selects one antenna  $(i_{s,p}, j_{s,p}) \in \mathcal{C}_{s,p}$  at each local snapshot. Hence, the local time index is an affine function of the selected antenna coordinates,  $p = a_s i_{s,p} + b_s j_{s,p} + c_s$ . The slopes  $(a_s, b_s)$  determine the contour orientation, and  $c_s$  shifts the intercept range  $p - c_s$  so that the selected local indices correspond to feasible contours of the finite array. Fig. 1 illustrates two affine sweeps on a small URA. Define  $y_{s,p} \triangleq y_{m_s^{(0)}+p}$ , and  $q_{s,p} \triangleq q_{m_s^{(0)}+p}$ . Substituting (5) and (6) in (4), the sample in sweep  $s$  becomes

$$y_{s,p} = \sum_{\ell=0}^{L-1} \alpha_{\ell,s} e^{j[(\omega_{x,\ell} + a_s \gamma_\ell) i_{s,p} + (\omega_{y,\ell} + b_s \gamma_\ell) j_{s,p}]} + q_{s,p},$$

where  $\alpha_{\ell,s} = \alpha_\ell e^{j\gamma_\ell(m_s^{(0)}+c_s)}$  is constant over  $p$ . Therefore, within sweep  $s$ , path  $\ell$  has the shifted spatial frequencies

$$\theta_{x,\ell,s} = \omega_{x,\ell} + a_s \gamma_\ell, \quad (7) \quad \theta_{y,\ell,s} = \omega_{y,\ell} + b_s \gamma_\ell. \quad (8)$$

The affine rule converts the temporal Doppler phase into a slope-dependent shift of the two-dimensional spatial spectrum. The Doppler parameter is therefore inferred from the change of the shifted spatial frequencies across sweeps. Let  $\boldsymbol{\eta}_\ell = [\omega_{x,\ell}, \omega_{y,\ell}, \gamma_\ell]^T$  and  $\boldsymbol{\theta}_\ell = [\theta_{x,\ell,1}, \theta_{y,\ell,1}, \dots, \theta_{x,\ell,S}, \theta_{y,\ell,S}]^T$ . Before modulo- $2\pi$  wrapping, (7) & (8) give the linear relation

$$\boldsymbol{\theta}_\ell = \mathbf{A}_S \boldsymbol{\eta}_\ell, \quad \mathbf{A}_S = \begin{bmatrix} 1 & 0 & \dots & 1 & 0 \\ 0 & 1 & \dots & 0 & 1 \\ a_1 & b_1 & \dots & a_S & b_S \end{bmatrix}^T. \quad (9)$$

The physical parameter vector  $\eta_\ell$  is identifiable from unwrapped shifted frequencies if  $\text{rank}(\mathbf{A}_S) = 3$ . This requires  $S \geq 2$  and at least two distinct slope pairs; for  $S = 2$ , it reduces to  $(a_1, b_1) \neq (a_2, b_2)$ . In practice, the shifted frequencies are observed modulo  $2\pi$ . We restrict attention to cases admitting a unique wrapping branch consistent with the assumed spatial and Doppler supports; otherwise, integer ambiguity resolution is required and is outside the present scope. The slope constraints below reduce the occurrence of such ambiguities.

### B. Slope Selection in Affine Sweeps

The slope pair  $(a_s, b_s)$  determines how Doppler is mapped into the shifted spatial frequencies. For sweeps  $r$  and  $s$ , the difference between the shifted frequencies of path  $\ell$  satisfies  $[\theta_{x,\ell,s} - \theta_{x,\ell,r} \quad \theta_{y,\ell,s} - \theta_{y,\ell,r}]^T = \gamma_\ell [a_s - a_r \quad b_s - b_r]^T$ . Thus, larger slope differences increase the sensitivity of the shifted spatial frequencies to Doppler. At the same time, large slopes also increase the Doppler-induced displacement  $a_s \gamma_\ell$  and  $b_s \gamma_\ell$ , which can cause modulo- $2\pi$  wrapping. We restrict the single-sweep Doppler displacement. For  $|\gamma| \leq \gamma_{\max}$ ,

$$|a_s| \gamma_{\max} \leq \rho_{\text{wrap}} \pi, \quad |b_s| \gamma_{\max} \leq \rho_{\text{wrap}} \pi, \quad (10)$$

where  $0 < \rho_{\text{wrap}} < 1$  is a margin to  $\pi$ -scale wrapping boundary. The slopes must also lead to feasible affine contours on the finite array. For a slope  $(a, b)$  and offset  $c$ , define the contour cardinality  $n_{a,b,c}(p) = |\{(i, j) \in \mathcal{I} : ai + bj + c = p\}|$ . A contour with  $n_{a,b,c}(p) = 0$  is infeasible, while a contour with very small cardinality gives little freedom to choose the antenna location on that contour. We therefore require that, for each selected slope, there exists an offset  $c$  such that

$$\min_{p \in \{0, \dots, M_s - 1\}} n_{a,b,c}(p) \geq n_{\min}, \quad (11)$$

where  $n_{\min} \geq 1$  is the minimum allowed contour occupancy. Large slopes tend to produce short or isolated contours on the finite array, reducing  $n_{a,b,c}(p)$ . Such contours make the antenna choice almost fixed and limit the role of the spatial prior in Sec. III-C. The candidate slope set is then

$$\begin{aligned} \mathcal{B} = \{ & (a, b) \in \mathbb{Z}^2 : (a, b) \neq (0, 0), \gcd(|a|, |b|) = 1, \\ & |a|, |b| \leq \lfloor \rho_{\text{wrap}} \pi / \gamma_{\max} \rfloor, \\ & \exists c \in \mathbb{Z} \text{ s.t. } \min_p n_{a,b,c}(p) \geq n_{\min} \}. \end{aligned} \quad (12)$$

The coprimality condition removes redundant contour orientations, since integer multiples such as  $(1, 2)$  and  $(2, 4)$  define the same family of parallel contours. After the slopes are selected, the corresponding offset for sweep  $s$  is chosen as

$$c_s \in \arg \max_{c \in \mathbb{Z}} \min_p n_{a_s, b_s, c}(p), \quad \text{s.t. (11)}. \quad (13)$$

Thus,  $c_s$  does not affect the Doppler-to-spatial shift in (7)–(8).

For a slope set  $\mathcal{A} = \{(a_s, b_s)\}_{s=1}^S \subset \mathcal{B}$ , we enforce  $\text{rank}(\mathbf{A}_S) = 3$  as required by (9). For  $L > 1$ , peaks generated by the same physical path must additionally be associated across sweeps before inversion.

The remaining multi-target bottleneck is the resolvability and association of the per-sweep shifted peaks. We therefore select slopes by testing whether nearby angle–Doppler paths can become unresolved in the shifted spatial domain after the affine Doppler-induced shift. Let  $d_\omega$

denote a spatial separation in discrete Fourier transform (DFT)-bin units and  $\phi$  its direction. We test  $\Delta\omega(d_\omega, \phi) = 2\pi [d_\omega \cos \phi / N_x \quad d_\omega \sin \phi / N_y]^T$ , where  $d_\omega \in [d_{\min}, d_{\max}]$ , and  $d_{\min}$  is the smallest targeted separation and  $d_{\max}$  limits the test to nearby paths for which aliasing is relevant. The Doppler-difference design set is  $\mathcal{G}_\Delta = \{\Delta\gamma : |\Delta\gamma| \leq 2\pi \Delta\nu_{\max} T\}$ , where  $\Delta\nu_{\max}$  is the maximum Doppler separation used in the design. For each  $(d_\omega, \phi, \Delta\gamma)$ , the shifted separation in sweep  $s$  is

$$\Delta\theta_s = \text{wrap}_{2\pi} \left( \Delta\omega(d_\omega, \phi) + \Delta\gamma [a_s \quad b_s]^T \right). \quad (14)$$

Small  $\|\Delta\theta_s\|_2$  means that two distinct angle–Doppler paths become close in the shifted spatial spectrum of sweep  $s$ , causing a per-sweep resolution or association ambiguity. We score  $\mathcal{A}$  using the lower-tail shifted-separation margin

$$d_{\text{sep}}(\mathcal{A}) = \text{perc}_{q_{\text{sep}}} \left\{ \min_s \|\Delta\theta_s\|_2 : d_\omega \in [d_{\min}, d_{\max}], \phi \in [0, 2\pi), \Delta\gamma \in \mathcal{G}_\Delta \right\}. \quad (15)$$

The percentile  $q_{\text{sep}}$  penalizes slope sets with frequent shifted-domain near-aliasing, favoring resolvable per-sweep peaks and stable cross-sweep association. The slope set is selected by

$$\begin{aligned} \max_{\mathcal{A} \subset \mathcal{B}, |\mathcal{A}|=S} d_{\text{sep}}(\mathcal{A}) \\ \text{s.t. } \text{rank}(\mathbf{A}_S) = 3, \quad \kappa(\mathbf{A}_S^T \mathbf{A}_S) \leq \kappa_{\max}. \end{aligned} \quad (16)$$

where  $\kappa(\cdot)$  denotes the spectral condition number and  $\kappa_{\max}$  limits noise amplification in the linear recovery.

Since  $\mathcal{B}$  contains only bounded integer slope pairs, the search is finite and performed offline. It is rerun only when the assumed Doppler support, array/sampling configuration, or resolution requirements change. The design parameters are set from the sensing scenario, estimator requirements, and finite-array geometry. With carrier wavelength  $\lambda_c$ , the maximum path-length rate  $\dot{\ell}_{\max}$  sets  $\gamma_{\max} = 2\pi T \dot{\ell}_{\max} / \lambda_c$ , while the maximum relative path-length rate between simultaneous targets sets  $\Delta\nu_{\max} = \Delta\dot{\ell}_{\max} / \lambda_c$ . If no pairwise prior is available, we use the conservative value  $\Delta\nu_{\max} = 2\dot{\ell}_{\max} / \lambda_c$ . The wrapping margin  $\rho_{\text{wrap}}$  leaves at least one grid/refinement cell to the  $\pi$ -scale wrapping boundary and determines the allowable slopes through (10). We use the smallest number of sweeps satisfying  $\text{rank}(\mathbf{A}_S) = 3$  and the per-sweep sample requirement of the 2D peak estimator; hence,  $S = 2$  is the default unless additional cross-sweep consistency is needed. The separation range in (15) follows the peak detector:  $d_{\min}$  is the non-overlap guard, and  $d_{\max}$  is the largest separation that can be shifted into near collision by the maximum expected Doppler difference. The percentile  $q_{\text{sep}}$  sets the tolerated lower-tail fraction of near-collision cases. Finally,  $n_{\min} = \max\{1, \lceil \beta_c N / M_s \rceil\}$ , with  $0 < \beta_c < 1$ , is chosen so that each contour retains enough antenna choices for aperture and coherence control.

### C. Contour-Wise Antenna Selection

For fixed slopes and offsets, the affine rule specifies the feasible contour  $\mathcal{C}_{s,p}$  at each local index  $p$ , but it does not uniquely specify which antenna on that contour should be selected. A uniform choice over  $\mathcal{C}_{s,p}$  is possible. However,

since the selected antennas determine the effective receive aperture and sidelobe behavior, we use an offline spatial prior to guide the contour-wise selection.

Let  $q_{ij}$  be a desired spatial sampling prior, with

$$q_{ij} \geq 0, \quad \sum_{i=0}^{N_x-1} \sum_{j=0}^{N_y-1} q_{ij} = 1. \quad (17)$$

At index  $p$  of sweep  $s$ , this prior is restricted to the feasible contour  $\mathcal{C}_{s,p}$ . The selected antenna is drawn according to

$$\Pr\{(i, j) \mid s, p\} = \frac{q_{ij}}{\sum_{(u,v) \in \mathcal{C}_{s,p}} q_{uv}}, \quad (i, j) \in \mathcal{C}_{s,p}. \quad (18)$$

The denominator is the prior mass on the feasible contour and only normalizes the prior over the antennas that satisfy the affine rule. The affine rule preserves the Doppler-aware space-time structure, while the prior determines which feasible antenna locations are preferred. For the prior design, define the selection intensity  $z_{ij} = Mq_{ij}$ , so that  $\sum_{i,j} z_{ij} = M$ . Let  $\tilde{\mathbf{u}}_{ij} = [i - (N_x - 1)/2, j - (N_y - 1)/2]^T$  be the centered antenna coordinate. Motivated by the array moment-of-inertia matrix [11], we define the weighted spatial aperture matrix

$$\mathbf{Q}_{\text{sp}}(\mathbf{Z}) = \sum_{i=0}^{N_x-1} \sum_{j=0}^{N_y-1} z_{ij} \tilde{\mathbf{u}}_{ij} \tilde{\mathbf{u}}_{ij}^T, \quad (19)$$

induced by  $\mathbf{Z} = [z_{ij}]$ . We use the scalar aperture score  $\log \det(\mathbf{Q}_{\text{sp}}(\mathbf{Z}) + \epsilon \mathbf{I}_2)$ , with regularization  $\epsilon > 0$ , since it promotes a profile that is spread in both array dimensions. Large aperture improves angular sensitivity [12], while spatial undersampling can create large Fourier sidelobes and grating-lobe ambiguities [13]. We thus use the sampling coherence

$$\mu(\mathbf{Z}) = \frac{1}{M} \max_{(\Delta k_x, \Delta k_y) \neq (0,0)} \left| \sum_{i,j} z_{ij} e^{j2\pi \left( \frac{\Delta k_x i}{N_x} + \frac{\Delta k_y j}{N_y} \right)} \right|$$

as a DFT-grid ambiguity metric. The prior is obtained from

$$\min_{\mathbf{Z}, t} \quad t \quad (20a)$$

$$\text{s.t.} \quad \mu(\mathbf{Z}) \leq t, \quad (20b)$$

$$\log \det(\mathbf{Q}_{\text{sp}}(\mathbf{Z}) + \epsilon \mathbf{I}_2) \geq \Gamma, \quad (20c)$$

$$\sum_{i=0}^{N_x-1} \sum_{j=0}^{N_y-1} z_{ij} = M, \quad (20d)$$

$$0 \leq z_{ij} \leq 1, \quad 0 \leq i < N_x, \quad 0 \leq j < N_y, \quad (20e)$$

for a target aperture level  $\Gamma$ . The objective limits the largest expected Fourier sidelobe, while the log-determinant constraint enforces a minimum spatial aperture. The problem is convex, since  $\mathbf{Q}_{\text{sp}}(\mathbf{Z})$  is affine in  $\mathbf{Z}$ ,  $\mu(\mathbf{Z})$  is convex, and the superlevel set of the concave log-determinant is convex [14]. After solving (20), we set  $q_{ij} = z_{ij}/M$  and generate the switching schedule offline via (18). The schedule is fixed during operation and adds no online computational overhead.

#### D. Affine Two-Stage Angle-Doppler Estimator

The proposed two-stage estimator first estimates the shifted spatial frequencies in each affine sweep and then associates these estimates across sweeps to recover the underlying angle-Doppler parameters. For sweep  $s$ , let  $\mathbf{y}_s = [y_m]$ , where  $m \in \mathcal{M}_s$ . For  $\boldsymbol{\theta} = [\theta_x, \theta_y]^T$ , define  $\mathbf{b}_s(\boldsymbol{\theta}) = [e^{j(\theta_x i_m + \theta_y j_m)}]$ . The coarse shifted-frequency estimates are obtained from the concentrated least-squares spectrum

$$P_s(\boldsymbol{\theta}) = |\mathbf{b}_s^H(\boldsymbol{\theta}) \mathbf{y}_s|^2 / \|\mathbf{b}_s(\boldsymbol{\theta})\|_2^2. \quad (21)$$

The  $L$  strongest non-overlapping peaks of  $P_s(\boldsymbol{\theta})$  are then refined locally using separable nonlinear least squares [15, Secs. 4.3], yielding the estimates  $\{\hat{\boldsymbol{\theta}}_{\ell,s}\}_{\ell=0}^{L-1}$ . Here,  $L$  is assumed known to the peak selector; alternatively, it can be estimated using a standard model-order criterion [16]. For  $L = 1$ , the cross-sweep association is unique. For  $L > 1$ , sweep 1 is taken as the reference and a candidate association is specified by  $\chi = (\pi_2, \dots, \pi_S)$ , where each  $\pi_s$  is a permutation of  $\{0, \dots, L-1\}$ . Under association  $\chi$ , the estimate  $\hat{\boldsymbol{\theta}}_{\pi_s(\ell),s}$  is assigned to the same path as  $\hat{\boldsymbol{\theta}}_{\ell,1}$ . The stacked shifted-frequency vector corresponding to path label  $\ell$  is

$$\hat{\boldsymbol{\theta}}_{\ell}^{(\chi)} = [\hat{\boldsymbol{\theta}}_{\ell,1}^T \quad \hat{\boldsymbol{\theta}}_{\pi_2(\ell),2}^T \quad \dots \quad \hat{\boldsymbol{\theta}}_{\pi_S(\ell),S}^T]^T. \quad (22)$$

For each candidate association, the affine-consistent angle-Doppler estimate is obtained by least squares (see (9)),

$$\hat{\boldsymbol{\eta}}_{\ell}^{(\chi)} = (\mathbf{A}_S^T \mathbf{A}_S)^{-1} \mathbf{A}_S^T \hat{\boldsymbol{\theta}}_{\ell}^{(\chi)}. \quad (23)$$

The selected association minimizes wrapped affine residual,

$$\hat{\chi} = \arg \min_{\chi \in \Pi} \sum_{\ell=0}^{L-1} \left\| \text{wrap}_{2\pi} \left( \hat{\boldsymbol{\theta}}_{\ell}^{(\chi)} - \mathbf{A}_S \hat{\boldsymbol{\eta}}_{\ell}^{(\chi)} \right) \right\|_2^2, \quad (24)$$

where  $\Pi$  denotes the set of all candidate associations and  $|\Pi| = (L!)^{S-1}$ . For  $S = 2$ , this reduces to a two-dimensional assignment problem. The final estimate is  $\hat{\boldsymbol{\eta}}_{\ell} = \hat{\boldsymbol{\eta}}_{\ell}^{(\hat{\chi})}$ . For peak detection, a minimum separation distance  $d_{\min}^{\theta}$  is enforced between selected peaks in each sweep. Specifically, the detected peaks satisfy  $\min_{\ell \neq r} \left\| \text{wrap}_{2\pi} \left( \hat{\boldsymbol{\theta}}_{\ell,s} - \hat{\boldsymbol{\theta}}_{r,s} \right) \right\|_2 \geq d_{\min}^{\theta}$ , where  $d_{\min}^{\theta}$  is the non-overlap distance used by the peak detector.

We summarize the online computational cost. Let  $G_{\theta} = G_x G_y$  denote the number of 2D spatial grid points,  $G_{\gamma}$  the number of Doppler grid points in a direct 3D search, and  $I_{\text{ref}}$  the number of local refinement iterations. The affine estimator requires  $\mathcal{O}(MG_{\theta})$  operations for the per-sweep coarse searches and  $\mathcal{O}(I_{\text{ref}} LM)$  operations for local refinement. For  $S = 2$ , association is a bipartite assignment problem with  $\mathcal{O}(L^3)$  complexity using the Hungarian algorithm [17]. For  $S > 2$ , exhaustive association scales as  $(L!)^{S-1}$  and can become the dominant cost. Hence, for the default  $S = 2$  case,

$$\mathcal{C}_{\text{aff}} = \mathcal{O}(MG_{\theta} + I_{\text{ref}} LM + L^3). \quad (25)$$

In comparison, a direct coarse angle-Doppler search requires  $\mathcal{O}(MG_{\theta} G_{\gamma})$  operations. Thus, the proposed estimator removes the Doppler grid from the dominant online search.

#### IV. NUMERICAL RESULTS

We evaluate the prior design, single-target angle-Doppler accuracy, two-target shifted-domain aliasing, and runtime scaling. Unless stated otherwise, the array is a  $32 \times 32$  half-wavelength URA with  $M = 250$  switched samples,  $S = 2$  sweeps, and  $T = 8.92 \mu\text{s}$ , so  $\gamma = 2\pi\nu T$ . We use the Doppler support  $\mathcal{D}_{\nu} = [-6, 6]$  kHz, corresponding to  $|\gamma| \leq 0.336$  rad per switching interval. In path-length rate, this is  $|\dot{\ell}| = \lambda_c |\nu|$ ; for  $\lambda_c \approx 1.1$  cm, the edge of  $\mathcal{D}_{\nu}$  corresponds to about 66 m/s, or half this radial velocity for a two-way monostatic echo. Slopes are selected using (16) with  $|a|, |b| \leq 6$ ,  $d_{\omega} \in [0.5, 4]$ ,  $\phi \in [0, 2\pi]$ ,  $\Delta\nu \in [-14, 0]$  kHz, and  $q_{\text{sep}} = 10$ . Spatial frequencies are reported in DFT-bin units,  $k_x = N_x \omega_x / (2\pi)$  and  $k_y = N_y \omega_y / (2\pi)$ . Prior design and




Fig. 2. Aperture-coherence tradeoff induced by the spatial prior design. Increasing the required aperture level  $\Gamma$  forces the prior toward larger array radii, improving angular aperture but increasing the largest Fourier sidelobe.

slope selection are performed once offline. For the single-target Cramér-Rao bound (CRB), with  $\mathbf{R} = [\mathbf{r}_0, \dots, \mathbf{r}_{M-1}]^T$  and  $\mathbf{P}_\perp = \mathbf{I}_M - \mathbf{1}_M \mathbf{1}_M^T / M$ , where  $\mathbf{I}_M$  and  $\mathbf{1}_M$  denote the identity matrix and all-ones vector, respectively, the equivalent Fisher information matrix (EFIM), after eliminating the unknown complex gain, is  $\mathbf{J}_\eta^e = 2|\alpha|^2 \mathbf{R}^T \mathbf{P}_\perp \mathbf{R} / \sigma^2$ . The Doppler CRB is obtained by converting from  $\gamma$  to  $\nu$ . The non-affine baseline uses a finite-grid 3D coarse search followed by concentrated nonlinear least-squares refinement. In the single-target study, this grid has  $33 \times 33 \times 81$  points over  $\mathcal{D}_\nu$ , while the affine estimator uses one  $49 \times 49$  shifted-spatial grid per sweep. In the two-target study, each affine sweep uses an  $81 \times 81$  grid, a 0.9-bin peak guard, and a 1.25-bin refinement window. Runtime tests use full DFT spatial grids and  $N_\gamma = 128$  Doppler points. The single-target case uses  $(k_x, k_y) = (7.52, -4.96)$  and  $\nu = 1200$  Hz. The two-target case fixes the first target at  $(-4, -3)$ , places the second at  $d_{\text{phys}} \in [2, 10]$  DFT bins along the Doppler-shift direction, and sweeps  $\Delta\nu$  across the predicted aliasing point. The signal-to-noise ratio (SNR) is the average received signal power over  $\sigma^2$ ; for one path,  $\text{SNR} = |\alpha|^2 / \sigma^2$ . The prior, single-target, two-target, and runtime studies use 300, 200, 50, and 100 Monte Carlo trials, respectively.

### A. Prior-Guided Sampling Profile

We first examine the spatial prior before imposing affine feasibility. Since  $z_{ij}$  in (20) is an expected selection intensity,  $\mu(\mathbf{Z})$  is an expected sampling-ambiguity metric rather than the coherence of one realized schedule. Fig. 2 shows that small  $\Gamma$  gives a nearly uniform prior with low expected sidelobes, while larger  $\Gamma$  moves probability mass toward the array boundary and increases aperture at the cost of larger sidelobes. We use  $\Gamma = 20.63$ , where the optimized-profile coherence becomes comparable to that of finite random masks drawn from the same prior. Table I shows that affine conditioning introduces no empty contours and keeps the finite-realization coherence in the same range. The aperture loss is small for the small-slope family and larger for the Doppler-sensitive slope family.

### B. Single-Target Angle-Doppler Accuracy

We next isolate the single-target case to separate acquisition threshold from high-SNR accuracy. We compare two affine slope families,  $(2, 3), (3, 2)$ , which induce smaller Doppler shifts, and  $(3, 5), (5, 2)$ , which provide stronger Doppler sensitivity. The non-affine baseline uses a coarse 3D search over

TABLE I

EFFECT OF AFFINE FEASIBILITY ON FINITE SAMPLING PROFILES. THE SUBSCRIPTS u AND a DENOTE UNCONSTRAINED AND AFFINE FINITE MASKS. COHERENCE IS IN DB, AND  $\mathcal{J}_{\text{ap}} = \log \text{DET}(\mathbf{Q}_{\text{sp}} + \epsilon \mathbf{I}_2)$ .

| $\Gamma$ | Slopes         | $\mu_u$ | $\mu_a$ | $\mathcal{J}_{\text{ap},u}$ | $\mathcal{J}_{\text{ap},a}$ | Empty |
|----------|----------------|---------|---------|-----------------------------|-----------------------------|-------|
| 19.93    | (2, 3), (3, 2) | -16.77  | -15.18  | 19.93                       | 19.96                       | 0     |
| 19.93    | (3, 5), (5, 2) | -16.81  | -14.85  | 19.93                       | 19.36                       | 0     |
| 20.55    | (2, 3), (3, 2) | -11.58  | -12.92  | 20.39                       | 20.25                       | 0     |
| 20.55    | (3, 5), (5, 2) | -11.64  | -10.77  | 20.39                       | 19.69                       | 0     |
| 21.17    | (2, 3), (3, 2) | -4.38   | -6.28   | 21.17                       | 20.86                       | 0     |
| 21.17    | (3, 5), (5, 2) | -4.38   | -5.42   | 21.17                       | 20.59                       | 0     |




Fig. 3. Single-target estimation versus SNR for  $\Gamma = 20.63$ .  $\bullet$ : non-affine 3D search;  $\blacklozenge$ : affine two-stage with small slopes  $(2, 3), (3, 2)$ ;  $\blacktriangledown$ : affine two-stage with large slopes  $(3, 5), (5, 2)$ ; black dashed —: CRB.

$(\omega_x, \omega_y, \gamma)$  followed by local concentrated nonlinear least-squares refinement. The affine estimator instead uses one 2D shifted-spatial search per sweep and recovers  $(\omega_x, \omega_y, \gamma)$  through (9). Fig. 3 shows the spatial-frequency and Doppler root-mean-square error (RMSE) versus SNR for  $\Gamma = 20.63$ . The main trend is the acquisition threshold. Once correctly initialized, the non-affine 3D search gives the best high-SNR accuracy, but its larger search space creates more noise-supported competing peaks at low SNR. The affine estimator reaches the high-accuracy region earlier because each coarse search is only two-dimensional and Doppler appears as a structured spatial shift. This behavior reflects the finite-grid initialization stage, where the 3D search has a larger set of candidate noise peaks. At high SNR, the affine estimator has a modest accuracy loss because it estimates shifted spatial frequencies separately in each sweep and combines them by linear inversion, rather than coherently optimizing the full angle-Doppler likelihood. The Doppler RMSE follows the same trend. The larger-slope family gives lower Doppler RMSE over most SNRs due to larger sweep-to-sweep shifts, but these shifts also increase sensitivity to noise and multi-target shifted-domain aliasing, as studied next.

### C. Two-Target Resolution and Shifted-Domain Aliasing

We next characterize the multi-target limitation of the affine two-stage estimator. In each sweep, resolvability is governed by the shifted spatial domain in (14), not by the physical angular separation alone. A Doppler difference can therefore enlarge or shrink the apparent separation between two targets. Reliable multi-target recovery requires three practical conditions. The shifted peaks in every sweep must satisfy the detector guard  $d_{\text{min}}^\theta$ , the Doppler-induced shifts must remain within the wrapping limits imposed in the slope design, and the residual in (24) must have a unique minimizer over the candidate permutations. The following study isolates the first condition by examining the case where a Doppler difference reduces the




Fig. 4. Two-target resolution under Doppler-induced affine shifts. (a) Shifted-domain separation in the stressed sweep. (b) Spatial-frequency RMSE of the affine two-stage estimator, with  $10^{-3}$  added for logarithmic visualization. The red dashed and white solid lines mark the predicted aliasing condition and one-bin shifted-separation boundary.




Fig. 5. Measured grid-search runtime versus array side length. The fitted scaling laws are approximately  $N_x^{2.06}$  and  $N_x^{1.86}$ , respectively. The vertical offset reflects the removal of the Doppler grid in the affine two-stage estimator.

shifted-domain separation inside one sweep; it does not test all association, wrapping, or model-order failures. We vary both the physical separation and the Doppler difference  $\Delta\nu = \nu_2 - \nu_1$ . The second target is placed along the Doppler-shift direction of a stressed sweep, for which the predicted aliasing point satisfies approximately  $d_{\text{phys}} + (2\pi T \Delta\nu) \|\mathbf{v}_s\|_2 \approx 0$ , where  $\mathbf{v}_s = [N_x a_s / (2\pi) \quad N_y b_s / (2\pi)]^T$ , and  $d_{\text{phys}}$  is measured in DFT-bin units. Fig. 4(a) shows the resulting shifted-domain separation. The red dashed line marks the predicted aliasing locus, and the white lines mark the one-bin shifted-separation boundary. Fig. 4(b) shows the corresponding spatial-frequency RMSE. The error ridge follows the low-separation region, confirming that the dominant failure mode is predictable shifted-domain aliasing. Away from this region, the estimator remains stable.

#### D. Complexity Scaling

Finally, we verify the runtime savings. The non-affine baseline uses  $N_x N_y N_\gamma$  grid points, while the affine method uses  $S N_x N_y$ , giving a nominal coarse-search reduction of  $N_\gamma / S$ . Fig. 5 shows that both runtimes scale approximately quadratically with  $N_x$ , with fitted trends  $N_x^{2.06}$  and  $N_x^{1.86}$ . The key gain is the vertical offset, as removing the Doppler grid reduces the average grid-search time from 538.7 ms to 5.52 ms ( $\approx 98\times$ ), and the largest-array time from 980.6 ms to 9.75 ms ( $\approx 101\times$ ). Including refinement, the average total runtime drops from 550.0 ms to 19.4 ms ( $\approx 28\times$ ).

#### V. CONCLUSION

We proposed affine antenna switching for low-complexity angle–Doppler estimation in single-RF ISAC sensing. The affine sampling rule converts Doppler-induced phase evolution

into deterministic shifts of the 2D spatial spectrum, enabling per-sweep 2D estimation followed by linear recovery of DoA and Doppler parameters. A probabilistic spatial prior further preserves aperture while controlling sampling coherence under affine feasibility. Numerical results show improved initialization at low SNR and lower runtime relative to the implemented direct 3D search, with shifted-domain aliasing as the main degradation mode. The current study assumes delay-isolated narrowband bins, known model order, and small numbers of paths; extensions to dense scenes requiring stronger association, wrapping resolution, and slope design for closely spaced angle–Doppler components are left for future work.

#### ACKNOWLEDGMENT

This work was supported in part by the European Commission through the MIDAS-6G project under Grant EU/HE/101207620, and in part by the Swedish Research Council (VR) under Grants 2022-03007 and 2024-04390.

#### REFERENCES

- [1] H. Chen *et al.*, “Mismatch analysis and cooperative calibration of array beam patterns for ISAC systems,” *IEEE Transactions on Wireless Communications*, vol. 25, pp. 16 496–16 511, 2026.
- [2] N. J. Myers *et al.*, “Swift-Link: A compressive beam alignment algorithm for practical mmWave radars,” *IEEE Transactions on Signal Processing*, vol. 67, no. 4, pp. 1104–1119, 2019.
- [3] J. Bechter *et al.*, “Compensation of motion-induced phase errors in TDM MIMO radars,” *IEEE Microwave and Wireless Components Lett.*, vol. 27, no. 12, pp. 1164–1166, 2017.
- [4] L. Liu *et al.*, “DOA estimation for switch-element arrays based on sparse representation,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2024, pp. 8506–8510.
- [5] M. Haardt *et al.*, “Joint estimation of 2-D arrival angles, propagation delays, and Doppler frequencies in wireless communications,” in *Proc. IEEE Digital Signal Processing Workshop*, vol. 1, 1998, pp. 1–4.
- [6] X. Geng *et al.*, “Joint velocity and angle estimation for single-cycle TDM MIMO automotive radar,” *Digital Signal Processing*, vol. 153, p. 104617, 2024.
- [7] M. K. Ercan *et al.*, “RIS-aided NLoS monostatic multi-target sensing under angle-doppler coupling,” *IEEE Transactions on Vehicular Technology*, vol. 74, no. 12, pp. 19 141–19 158, 2025.
- [8] Q. Wang *et al.*, “A high-resolution joint angle-doppler estimation Sub-Nyquist radar approach based on matrix completion,” *Information*, vol. 10, no. 4, 2019.
- [9] Y. Zhang *et al.*, “Random time division multiplexing based MIMO radar processing with tensor completion approach,” *Sensors*, vol. 23, no. 10, 2023.
- [10] A. Simončič *et al.*, “Experimental evaluation of enhanced antenna switching for CFO mitigation in DoA estimation,” *IEEE Open Journal of Antennas and Propagation*, vol. 6, no. 4, pp. 1022–1036, 2025.
- [11] U. Baysal *et al.*, “On the geometry of isotropic arrays,” *IEEE Transactions on Signal Processing*, vol. 51, no. 6, pp. 1469–1478, 2003.
- [12] W. Ma *et al.*, “Movable antenna enhanced wireless sensing via antenna position optimization,” *IEEE Transactions on Wireless Communications*, vol. 23, no. 11, pp. 16 575–16 589, 2024.
- [13] A. Gupta *et al.*, “Design of large effective apertures for millimeter wave systems using a sparse array of subarrays,” *IEEE Transactions on Signal Processing*, vol. 67, no. 24, pp. 6483–6497, 2019.
- [14] S. Boyd *et al.*, *Convex optimization*. Cambridge university press, 2004.
- [15] P. Stoica *et al.*, *Spectral Analysis of Signals*. Upper Saddle River, NJ, USA: Prentice Hall, 2005.
- [16] M. Wax *et al.*, “Detection of signals by information theoretic criteria,” *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 33, no. 2, pp. 387–392, 1985.
- [17] D. F. Crouse, “On implementing 2d rectangular assignment algorithms,” *IEEE Transactions on Aerospace and Electronic Systems*, vol. 52, no. 4, pp. 1679–1696, 2016.