International Journal of Scientific Advances
ISSN: 2708-7972
Volume: 7 | Issue: 6 | Nov – Dec 2026 Available Online: www.ijscia.com
DOI: 10.51542/ijscia.v7i6.34

Quantization-Aware Coordinate Descent for Low-Bit Phased Arrays: Joint Optimisation of Peak Sidelobe Level, Beam-Pointing Error and Gain Loss

Md Imran Hossain[1*], Morium Akter Boishakhi[2], Rafin Sami Uzzaman[1], Abdullah Al Mamun Sadi[1], Mahmud Tanjoy[1]

[1]School of Electronics and Information Engineering, Nanjing University of Information Science and Technology, Nanjing, China

[2]School of Computer Science and Technology, Taizhou University, Taizhou, China

E-mail: hossainmdimran@nuist.edu.cn; moriumboishakhi756@gmail.com; 202453050028@nuist.edu.cn; abdullahalmamunsadi@nuist.edu.cn; mahmudtanjoy@nuist.edu.cn

*Corresponding author details: Md. Imran Hossain; hossainmdimran@nuist.edu.cn

ABSTRACT

Low-resolution digital phase shifters reduce phased-array cost and power, but their coarse phase states raise the sidelobes, create discrete quantization lobes, tilt the beam, and reduce the gain. Existing control-domain remedies usually address only one of these effects. This study proposes quantization-aware coordinate descent (QACD), which selects the phase-shifter states by minimising a scalarised cost of peak sidelobe level (PSL), gain loss and pointing error, using exact per-element enumeration, rank-one pattern updates and a smoothed-maximum continuation; no hardware change is needed. First-order laws for the expected gain loss and the root-mean-square (RMS) pointing error under any symmetric taper are restated and verified by Monte Carlo simulation. QACD is compared with direct rounding, random phase offsets, the gain-optimal discrete beamformer, and a genetic algorithm in array-factor simulations of a 32-element, −30 dB Taylor linear array scanned over 0°–60° and of a 16 × 16 planar array. With 3-bit phase shifters, QACD lowers the mean PSL from −16.7 to −21.0 dB and the worst-case PSL from −14.0 to −19.3 dB and reduces the RMS pointing error from 0.161° to 0.037°, at 0.14 dB additional gain loss relative to direct rounding; with fixed random pre-phases, the mean PSL reaches −22.4 dB. Over ten seeds, QACD matches the genetic algorithm in mean PSL with 2.1 times fewer cost evaluations and improves the planar-array PSL from −16.9 to −27.8 dB. The benefit grows with array size but vanishes when random hardware phase errors approach the quantization error. Results come from an ideal array-factor model without full-wave or hardware validation.

Keywords: Phased array; Phase quantization; Low-resolution phase shifter; Peak sidelobe level; Beam-pointing error; Coordinate descent

INTRODUCTION

Electronically steered phased arrays are now used far beyond military radar: in millimetre-wave cellular base stations and terminals, in satellite user terminals and in reconfigurable intelligent surfaces (RIS). In each of these systems, the beam is steered by setting the phase of every element, and that phase is usually produced by a digital phase shifter with B control bits and therefore only L=2^B states [1]. The number of bits is a direct cost driver. Each extra bit adds a switching stage with its own insertion loss, chip area, control line, and calibration effort, and in large arrays these costs are multiplied by hundreds or thousands of elements.

This pressure has made low-resolution phase control an active research topic. Hybrid analog–digital millimetre-wave transceivers use large numbers of low-resolution phase shifters to save power [2, 3]; low-cost satellite terminals face the same trade-off between resolution and beam quality [4], and RIS prototypes are usually built with 1-bit or 2-bit elements whose quantization lobes must be suppressed [5-8].

The price of coarse phase states is a degraded beam. The realised phases differ from the ideal progressive phase by up to half a least-significant bit (LSB), and these errors have three visible effects [1], [9].

They reduce the gain in the commanded direction. They raise the sidelobes and, because the errors produced by rounding a linear phase are periodic across the aperture, they can create discrete quantization lobes far above the designed sidelobe level. They also tilt the beam away from the commanded direction [10].

Two families of remedies exist (Table 1). Hardware remedies change the array. Fixed random pre-phases decorrelate the quantization error and remove the discrete lobes [9]; recent RIS designs implement them with delay lines or randomised unit cells, demonstrated on fabricated, measured surfaces [5], [7] and in full-wave simulation [8].

Structured prephases combined with very-low-resolution phase shifters lower the sidelobes and the pointing error at a small gain cost [11]; subarray division with mixed-resolution phase shifters [4], modified aperture shapes [6] and dense delta-sigma arrays [12] are further options. Control-domain remedies keep the hardware and choose the states more cleverly. The statistics of the pointing deviation caused by random phasing rules were analysed for tapered planar arrays [10]. Linear and integer linear programming have been used to minimise pointing error [13], [14], simulated annealing to improve sidelobes [15], linear-time combinatorial algorithms to maximise gain exactly [16], [17], and genetic algorithms (GA) for general phase-only synthesis [18].

TABLE 1: Related Work on Low-Resolution Phase Control, Newest First.

Ref.	Year	Setting	Metric targeted	Method	Hardware change
[8]	2026	1-bit hybrid RIS	Quantization lobe	Phase randomisation via slot sizes (full-wave simulation)	Yes
[7]	2025	1-bit mmWave RIS	Quantization lobe	Fixed random phase delays (measured)	Yes
[4]	2025	Satellite user terminal	Pointing, PSL	Mixed-resolution phase shifters	Yes
[18]	2024	Planar phased array	Pattern restoration	Phase-only GA	No
[11]	2023	32 × 32 phased array	PSL, pointing, gain	Prephases + hybrid phasing	Yes
[6]	2023	1-bit RIS	Quantization lobe	Modified aperture shape	Yes
[17]	2023	RIS	Gain (SNR)	Linear-time global optimum	No
[16]	2022	RIS	Gain (SNR)	Linear-time approximation	No
[3]	2021	Hybrid mmWave MIMO	Spectral efficiency	Lattice (Babai) design	No
[5]	2020	1-bit reflective surface	Quantization lobe	Random pre-phase (measured)	Yes
[14]	2019	Planar phased array	Pointing	Integer linear programming roundoff	No
[13]	2016	Linear phased array	Pointing	Linear programming	No
[15]	2015	Radar phased array	Sidelobes	Simulated-annealing bit assignment	No
[10]	2003	Tapered planar array	Pointing	Statistics of random phasing methods	Varies
[9]	1983	Phased array	Quantization lobes	Random phasing methods compared	Varies

Each control-domain method targets one figure of merit, whereas a designer needs a low peak sidelobe level (PSL), a small pointing error and a small gain loss at the same time. Optimising one can silently worsen the others, and published comparisons rarely use identical arrays, scan sets and metrics. In the literature located for this study, no low-complexity control-domain method optimises all three together with an explicit trade-off and is benchmarked against both the gain-optimal discrete solution and a metaheuristic on identical metrics.

The objective of this study is to fill that gap with a simple, scripted method. Coordinate descent itself is standard; the contribution lies in its specialisation to this problem and in the evaluation: (i) error laws for the gain loss and the pointing error, restated in a form that guides optimisation and verified by Monte Carlo simulation; (ii) quantization-aware coordinate descent (QACD), which minimises a composite cost of PSL, gain loss and pointing error by exact per-element enumeration; (iii) a head-to-head comparison with four baselines across 1–6 bits, including a ten-seed comparison with a GA; and (iv) sensitivity studies on the cost weights, array size, amplitude taper, random hardware error and an analytical mutual-coupling model. No antenna geometry is designed, and no full-wave solver is used: all results come from the array-factor model with isotropic elements. The study addresses array-control and signal-processing practice rather than antenna-hardware design. The joint optimisation is scalarised: the weights ρ and κ select one point on the trade-off between the three metrics, and Pareto optimality is not claimed.

MATERIALS AND METHODS

Hardware Context

Figure 1 shows the kind of hardware addressed: a 16 × 16 tile that stacks a radiating layer of patch elements at half-wavelength pitch, a beamformer layer of multi-channel B-bit phase-shifter ICs, a feed layer with a corporate network and optional fixed pre-phase lines, and a control layer with a microcontroller or FPGA.

Each element is driven through the chain of Figure 2. The proposed method lives entirely in the control path: QACD runs offline on the array-factor model, its states are stored in a beam table, and the controller writes the B-bit control words whenever the beam is moved. Both figures are conceptual illustrations; no hardware was designed, simulated electromagnetically, fabricated, or measured.

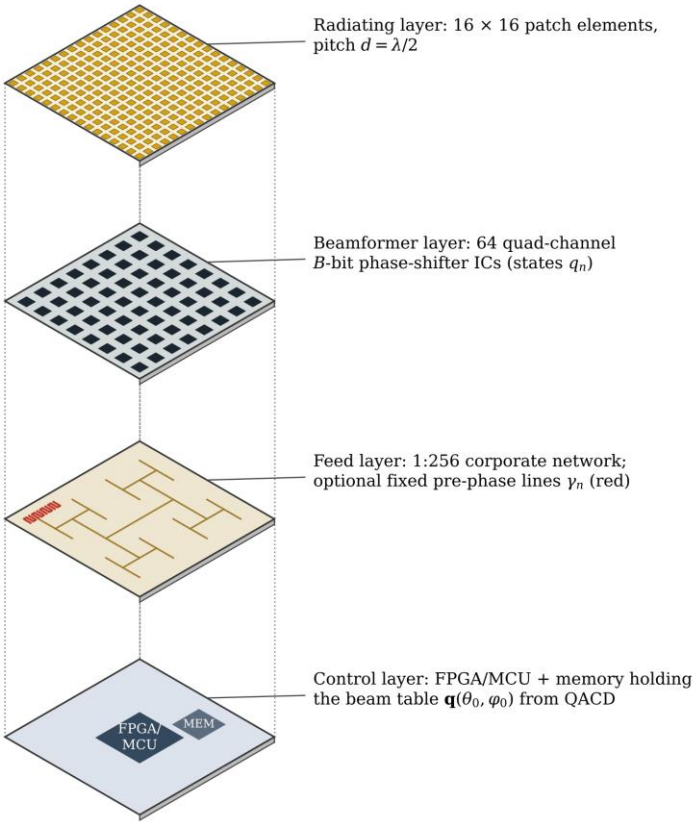

FIGURE 1: Conceptual Exploded View of a 16 × 16 Low-Bit Phased-Array Tile (Not to Scale). From Top: Radiating Layer, Beamformer Layer with B-Bit Phase-Shifter ICs, Feed Layer with the Corporate Network and Optional Pre-Phase Lines, and Control Layer Holding the Beam Table Computed by QACD.

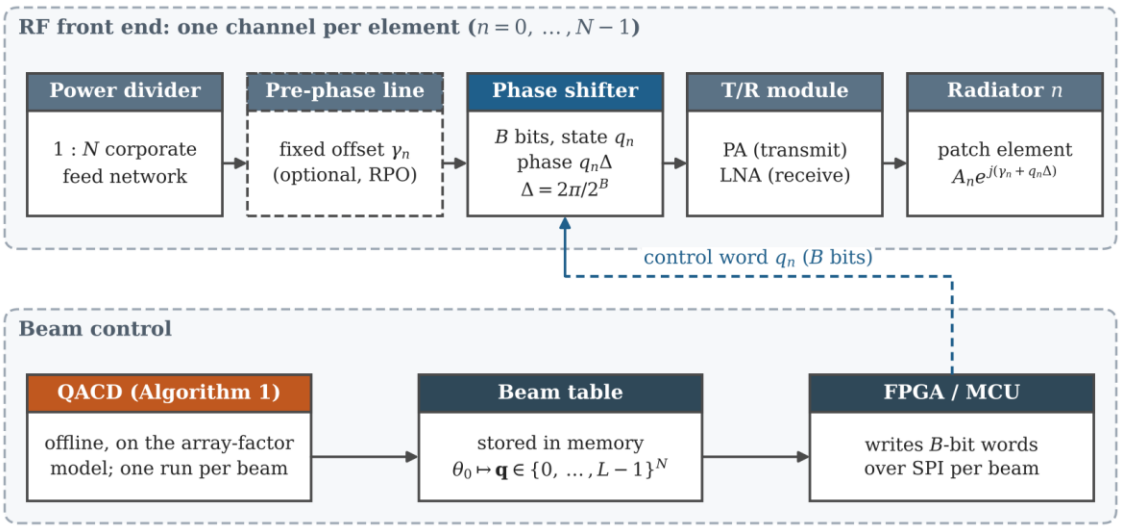

FIGURE 2: Signal Chain of Element n (Top) and the Beam-Control Path (Bottom). QACD Runs Offline and Fills the Beam Table; the Controller Writes the Control Word q_n to Each Phase Shifter. Illustration Only.

Array-Factor and Phase-Shifter Model

A narrowband, far-field linear array of N isotropic elements with spacing d is considered. Element n sits at $x_n = m_n d$ with centred index $m_n = n - (N-1)/2$. With $u = \sin\theta$, θ measured from broadside, and wavenumber $k = 2\pi/\lambda$, the array factor is

$$\mathrm{AF}(u) = \sum_{n=0}^{N-1} A_n \, e^{\,j(kd\, m_n u + \beta_n)} \qquad (1)$$

where $A_n > 0$ is a fixed symmetric amplitude taper ($A_n = A_{N-1-n}$) and β_n is the realised phase. The power pattern is $G(u) = |\mathrm{AF}(u)|^2$. Steering to $u_0 = \sin\theta_0$ requires the continuous phase

$$\beta_n^{\star} = -kd\, m_n u_0. \qquad (2)$$

A B-bit phase shifter realises one of $L = 2^B$ uniformly spaced states with least-significant bit $\Delta = 2\pi/L$:

$$\beta_n = \gamma_n + q_n \Delta, \qquad q_n \in \{0,1,\ldots,L-1\}. \qquad (3)$$

Here γ_n is a fixed, known pre-phase: zero for a conventional array, and drawn once uniformly over $[0, \Delta)$ for the random-phase-offset (RPO) feed [9]. The phase error is $\varepsilon_n = \mathrm{wrap}(\beta_n - \beta_n^{\star})$.

For the 16 × 16 uniform rectangular array the phase term becomes $kd(\tilde{x}_n u + \tilde{y}_n v)$ with centred indices $\tilde{x}_n, \tilde{y}_n$, $u = \sin\theta\cos\varphi$, $v = \sin\theta\sin\varphi$, and a separable taper. In the robustness study the radiated excitation is $\mathbf{Cw}$, where $\mathbf{w} = [A_n e^{j\beta_n}]$ and $\mathbf{C}$ is the banded symmetric Toeplitz coupling matrix widely used for uniform linear arrays [19], with $c_0 = 1$ and $c_l = \alpha^l e^{jl\psi}$ for $1 \le l \le P_c$. This coupling model is an analytical stress test with parameters chosen by us, not a prediction for a specific antenna. The isotropic array factor is a standard model [1]; it excludes element patterns and wideband squint, and the ideal uniform states exclude state-dependent amplitude and phase errors.

Performance Metrics

All metrics are computed on grids that are finer than, and independent of, the grids used inside the optimisers. The gain loss is $\mathrm{GL} = 10\log_{10}(|\mathrm{AF}^{\star}(u_0)|^2/|\mathrm{AF}(u_0)|^2)$ dB, relative to continuous phases; for isotropic elements at $d = \lambda/2$ it equals the directivity loss, because phase-only quantization leaves $\sum A_n^2$ unchanged. The PSL is the maximum of $G(u)$ over the sidelobe region $|u - u_0| \ge u_{\mathrm{null}}$ divided by the main-beam peak, where u_{null} is the first-null half-width of the continuous pattern, so quantization lobes always count as sidelobes. The integrated sidelobe level (ISL) is the sidelobe-region energy divided by the total radiated energy. The pointing error is $\delta u = u_{\mathrm{pk}} - u_0$, where u_{pk} is found on a fine local grid with three-point parabolic interpolation; in angle, $\Delta\theta = \arcsin(u_{\mathrm{pk}}) - \theta_0$. The half-power beamwidth (HPBW) is measured in θ around the realised peak.

Analytical Error Laws

Substituting $\beta_n = \beta_n^{\star} + \varepsilon_n$ into (1) gives, for any offset δ from the commanded direction,

$$\mathrm{AF}(u_0 + \delta) = \sum_n A_n \, e^{\,j(kd\, m_n \delta + \varepsilon_n)}, \qquad (4)$$

so, the shape of the beam is fully determined by the error sequence $\{\varepsilon_n\}$.

- *Status of the Laws*

Equation (5) is exact for independent errors uniform on $[-\Delta/2, \Delta/2)$. Equations (6)–(8) are first-order (small-argument) approximations of it and of the peak condition; they are not exact results for DR or QACD, whose errors are neither independent nor uniform. Their accuracy against Monte Carlo simulation is quantified in the Results section (Validation of the Error Laws).

- *Gain Loss*

If the ε_n are independent and uniform on $[-\Delta/2, \Delta/2)$, which holds exactly for the RPO feed and approximately for rounding averaged over many angles, then $\mathbb{E}[e^{j\varepsilon}] = \mathrm{sinc}(\Delta/2)$ with $\mathrm{sinc}(x) = \sin x/x$, and

$$\frac{\mathbb{E}\,|\mathrm{AF}_q(u_0)|^2}{|\mathrm{AF}^{\star}(u_0)|^2} = \mathrm{sinc}^2\left(\frac{\Delta}{2}\right) + \left[1 - \mathrm{sinc}^2\left(\frac{\Delta}{2}\right)\right]\eta, \qquad \eta = \frac{\sum_n A_n^2}{(\sum_n A_n)^2}. \qquad (5)$$

The first term is the classical gain-reduction factor [1], [9] and η is a finite-array correction ($\eta = 1/N$ for uniform amplitudes, 0.0365 for the Taylor taper used here). For small Δ,

$$\mathrm{GL} \approx 4.343\,(1-\eta)\,\Delta^2/12 \ \ \mathrm{dB}, \qquad (6)$$

so, the gain loss falls by a factor of about four per bit.

- *Pointing Error*

The realised peak satisfies $dG/d\delta = 0$. Linearising for small arguments and using the symmetric-taper identities $\sum A_n m_n = 0$ and $\sum_n \sum_{n'} A_n A_{n'} (m_n - m_{n'})^2 = 2S_0 S_2$, with $S_0 = \sum A_n$ and $S_2 = \sum A_n m_n^2$, gives

$$\delta u \approx -\frac{\sum_n A_n\, m_n\, \varepsilon_n}{kd\, S_2}. \qquad (7)$$

Thus δu is the amplitude-weighted least-squares slope of the phase error across the aperture: only the tilt component of the error moves the beam. Linear relations of this kind are classical [10], [13]; (7) restates the relation in the form that enters the cost function below. For independent errors with variance $\Delta^2/12$,

$$\delta u_{\mathrm{rms}} = \frac{\Delta}{\sqrt{12}} \frac{\sqrt{\sum_n A_n^2 m_n^2}}{kd \sum_n A_n m_n^2} \xrightarrow{A_n = 1,\ d = \lambda/2} \frac{2^{1-B}}{\sqrt{N(N^2-1)}} \qquad (8)$$

so, the RMS error halves with each bit and falls as $N^{-3/2}$.

- *Rounding and the 1-Bit Mirror Lobe*

With direct rounding (DR), ε_n is periodic in n whenever $kd\,u_0/\Delta$ is rational. An error sequence with period T radiates discrete lobes at $u = u_0 + 2\ell/T$ for $d = \lambda/2$, $\ell = \pm1, \pm2, \ldots$ [1], and a long period adds a linear component that (7) maps into a pointing error; both effects make DR worse than the independent-error model. For $B = 1$ and $\gamma_n = 0$ every weight is real, so $|\mathrm{AF}(-u_0)| = |\mathrm{AF}(u_0)|$: a 1-bit array without pre-phases always has a 0 dB mirror lobe, which no optimiser can remove.

Proposed Method: Quantization-Aware Coordinate Descent

QACD chooses the state vector $\mathbf{q} = (q_0, \ldots, q_{N-1})$ by minimizing

$$J(\mathbf{q}) = \mathrm{PSL}_{\mathrm{dB}}(\mathbf{q}) + \kappa\,\mathrm{GL}_{\mathrm{dB}}(\mathbf{q}) + \rho\left(\frac{\delta u(\mathbf{q})}{u_{\mathrm{tol}}}\right)^2, \qquad (9)$$

where $\kappa \geq 0$ prices one decibel of gain loss in decibels of PSL and $\rho \geq 0$ prices a pointing error equal to the tolerance u_{tol}. Because the PSL is a maximum, J is flat for most single-element moves. The maximum is therefore first replaced by the generalised p-mean of the sidelobe power samples over the sidelobe region Ω,

$$M_p = \max_{\Omega} G\left[\frac{1}{|\Omega|}\sum_{u\in\Omega}\left(\frac{G(u)}{\max_{\Omega} G}\right)^p\right]^{1/p} \xrightarrow{p\to\infty} \max_{\Omega} G, \qquad (10)$$

and QACD is run to convergence for $p = 4$, then 16, then ∞ (the true PSL), warm-starting each stage. A small p rewards moves that lower many sidelobes at once. Changing element n from state q_n to state s changes the pattern by one term,

$$\mathrm{AF}^{(s)}(u) = \mathrm{AF}(u) + A_n\left(e^{\,j(\gamma_n + s\Delta)} - e^{\,j(\gamma_n + q_n\Delta)}\right) e^{\,jkd\,m_n u}, \qquad (11)$$

so, with AF stored on the sidelobe grid (K points), a local peak grid (K_l points) and at u_0, all L states of one element are evaluated in $\mathcal{O}\big(L(K + K_l)\big)$ operations, and one sweep costs.

$$\mathcal{O}\big(NL(K + K_l)\big) \text{ operations and } \mathcal{O}\big(N(K + K_l)\big) \text{ memory.} \qquad (12)$$

The algorithm is initialised with the gain-optimal discrete beamformer (OGM), obtained by testing every arc of the common rotation phase [17] (with pre-phases, on the offset phases: RPO+QACD). In each stage the elements are visited in a random order with a fixed seed; for each element, all L states are evaluated with (11), and the best is kept if it lowers J strictly. A stage ends after a sweep with no change or after 50 sweeps. Because the current state is always a candidate, J never increases within a stage, and each stage ends at a local minimum where no single-element change lowers J; global optimality is not claimed. The true cost can rise at a stage switch, because the objective changes. QACD needs only the array model, so it runs unchanged with pre-phases, with a coupling matrix ($\mathbf{C}$ enters the steering vectors) or on a planar geometry.

Simulation Setup and Baselines

The main array is a 32-element, half-wavelength linear array with a −30 dB Taylor taper (Table 2). With continuous phases it keeps a −30.2 dB PSL at every scan angle, and its HPBW widens from 4.02° at broadside to 8.11° at 60°. The weights $\rho = \kappa = 1$ were fixed before any experiment and were not tuned; they are the neutral choice that prices one decibel of gain loss like one decibel of PSL and a pointing error equal to the tolerance like one decibel. The sensitivity study on the cost weights shows that this choice does not change any conclusion. QACD is compared with four baselines that share the same state table and metric code: DR, the default in most beam-steering controllers; RPO, rounding with fixed random pre-phases [9]; OGM, the exact gain-optimal discrete beamformer [17]; and an integer GA minimising the same cost (9), representing metaheuristic phase-only synthesis [18]. QACD-nc is an ablation without continuation (p=∞ only).

TABLE 2: Simulation Parameters.

Parameter	Value
Main array	Linear, $N = 32$, $d = \lambda/2$, isotropic elements; Taylor −30 dB, $\bar{n} = 5$ ($\eta = 0.0365$); uniform and Taylor −35 dB in Table 6
Phase-shifter bits	$B = 1, \ldots, 6$ (theory checked up to $B = 8$)
Scan angles	0°–60°, 1° step; 3° step in the weight, hardware-error and coupling studies; 4° step in the array-size study
Optimisation grids	2048 points in u; local peak grid ±0.75 HPBW, 101 points
Reporting grids	8192 points in u; local peak grid ±1.0 HPBW, 801 points
Sidelobe region	$\lvert u - u_0\rvert \geq u_{\mathrm{null}} = 0.0940$; planar array: outside the first-null square
QACD	$\rho = \kappa = 1$, $u_{\mathrm{tol}} = 0.02\,\mathrm{HPBW}_u = 1.40 \times 10^{-3}$ (0.080° at broadside); $p = 4 \to 16 \to \infty$; at most 50 sweeps per stage
GA	Population 40, 100 generations (3840 evaluations), binary tournament, uniform crossover 0.9, mutation $1/N$, elitism 2, seeded with OGM and DR
RPO	γ_n uniform on $[0, \Delta)$, 10 fixed realisations (seeds 100–109); RPO+QACD uses seeds 100–102
Planar array	16 × 16, separable Taylor −30 dB, $(\theta_0, \varphi_0) = (30°, 40°)$, $B = 2,3$
Robustness	Phase error 0–15° RMS; Toeplitz coupling $\alpha = 0$–0.3, $\psi = \pi/4$, $P_c = 2$; $N = 16$–128
Software and computer	Python 3.12.3, NumPy 2.5.4 [20], SciPy 1.18.1 [21], Matplotlib 3.11.2; one CPU core of a cloud container

Statistical Analysis

All experiments are scripted with fixed random seeds and are therefore repeatable. States are optimised on the coarse grids and evaluated on the finer, independent reporting grids, so no method benefits from fitting its own sampling grid. Scan-level results are reported as the mean and the worst case over the scan angles; RPO results are averaged over the offset realisations. The error laws are checked against Monte Carlo simulation with 20 000 independent error draws per bit number, the hardware-error study uses 200 random trials per error level, and the comparison between QACD and GA is repeated with ten seeds and reported as mean ± sample standard deviation with the range. Cost evaluations count evaluations of J for any candidate state vector, including the NL evaluations of the OGM initialisation for QACD. Run times are wall-clock times of unoptimised Python and are indicative only.

RESULTS AND DISCUSSION

Physical Origin of the Quantization Lobe

At θ_0=22°, the worst DR angle on the 1° grid, the progressive phase per element is 67.4°, almost exactly 1.5 LSB at B=3. Rounding therefore alternates between errors of about ±11° (Figure 3a), a period-2 sequence that radiates discrete lobes at u-u_0=±1 (Figure 3b). The visible one appears at θ=-38.7° with a level of −14.0 dB (Figure 4a): this is the quantization lobe. OGM produces the same pattern as DR here, because maximising gain does not care where the error energy goes. QACD breaks the periodicity and spreads the error energy into a nearly flat floor, reaching −20.4 dB PSL; GA reaches −19.3 dB and RPO+QACD −22.8 dB (Figure 4b). The price is a far-sidelobe floor between about −20 and −25 dB instead of the −30 dB Taylor level.

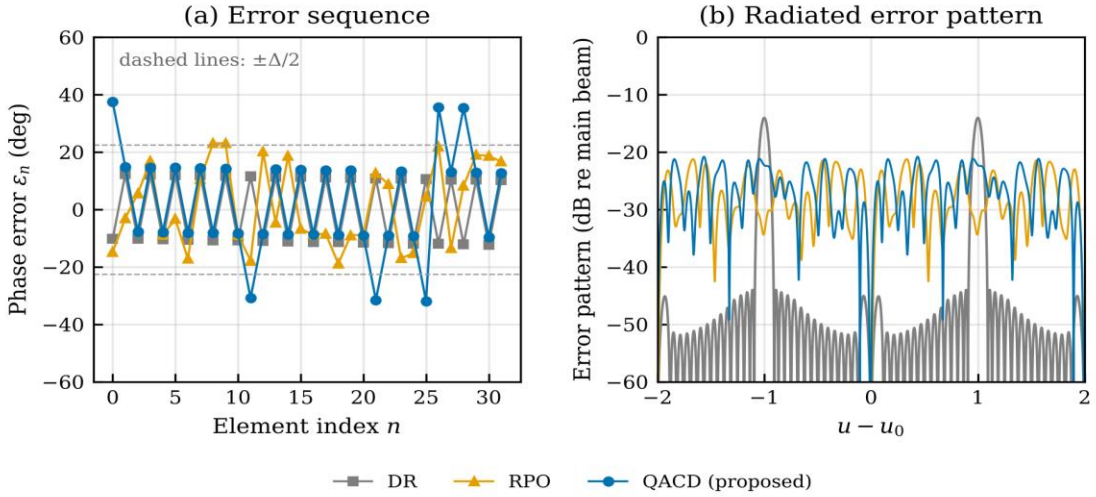

FIGURE 3: Phase-Error Sequences at B=3, θ_0=22° (Common Phase Removed) and the Patterns They Radiate. (a) DR Produces a Period-2 Error; RPO and QACD Do Not. (b) The DR Error Radiates Discrete Lobes at u-u_0=±1, Whereas the RPO and QACD Errors Spread Their Energy over Angle.

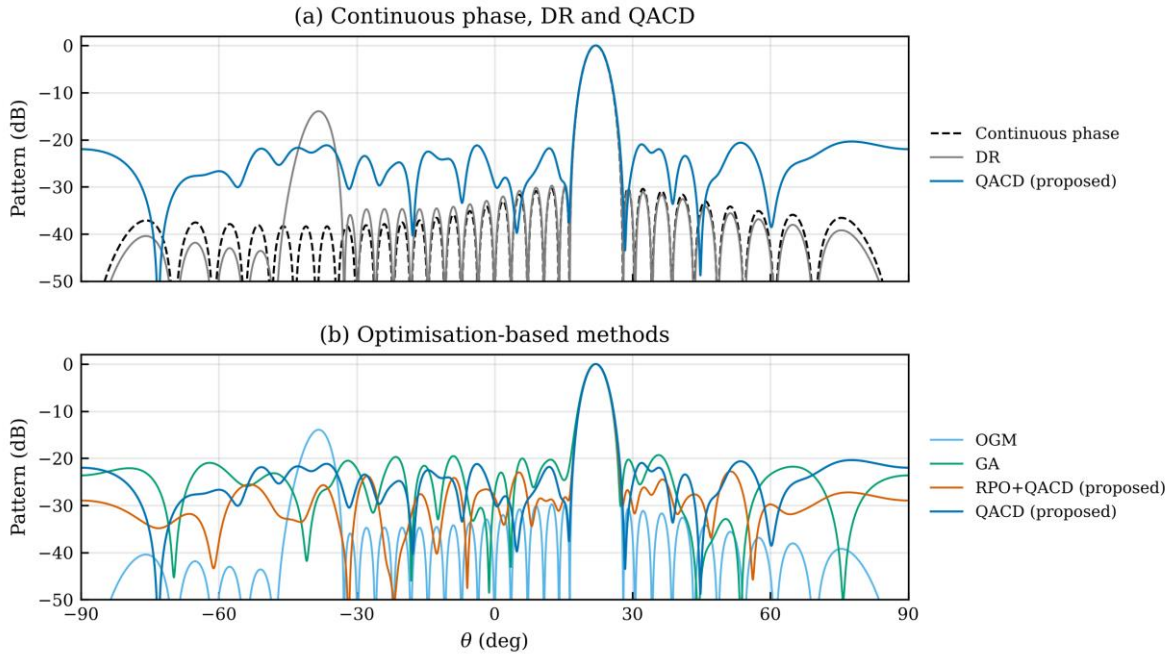

FIGURE 4: Radiation Patterns at B=3, θ_0=22°. (a) Continuous Phase, DR and QACD. (b) OGM, GA, RPO+QACD (One Offset Realisation) and QACD.

Peak Sidelobe Level Versus Resolution

For B=2–4, DR improves by 5.5–7.0 dB per bit in mean PSL and then saturates towards the −30 dB taper floor (Figure 5). QACD gains 3.9–4.3 dB in mean PSL over DR for B=2–4 and 5.3 dB in the worst case at B=3; at B=6 it reaches −30.1 dB, within 0.1 dB of the continuous design.

At B=1, every method without pre-phases has a 0 dB worst case, the mirror lobe derived in the error-law analysis; RPO reaches −6.7 dB mean PSL and RPO+QACD −7.5 dB. QACD has a better mean PSL than GA for B≥3, whereas GA is better at B=2 (−14.5 versus −13.6 dB). Per angle, GA reaches the lower PSL at 43% of the angles at B=3 and at 89% at B=2, so neither method dominates angle by angle.

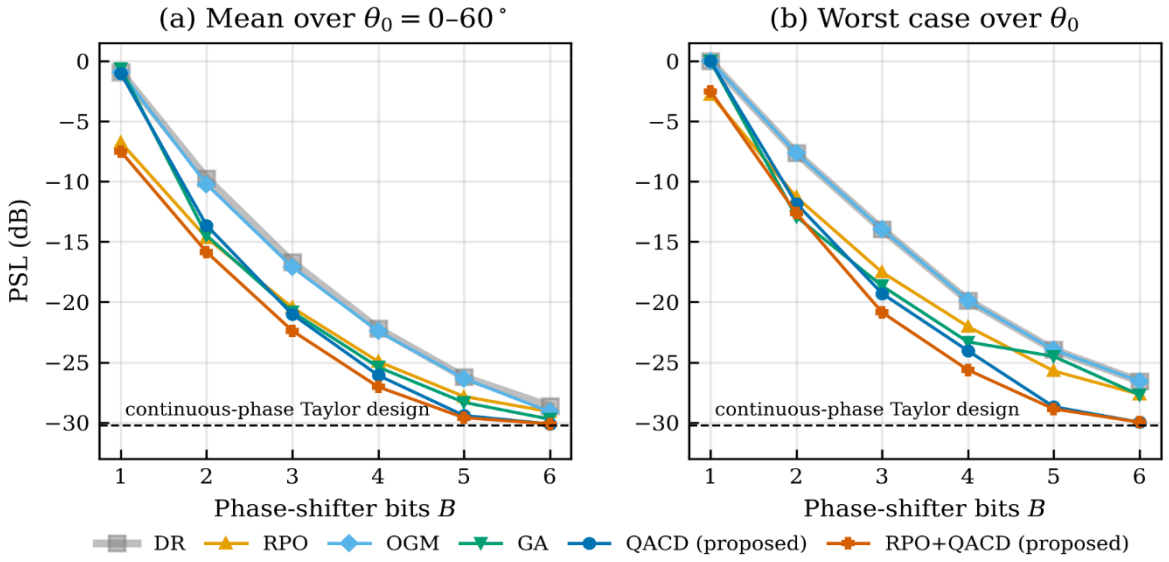

FIGURE 5: PSL Versus Number of Bits B. (a) Mean over θ_0=0°–60°. (b) Worst Case over θ_0 (RPO-Based Methods Averaged over Offset Realisations). DR Is Drawn Wide Because OGM lies on Top of it in (b).

Validation of the Error Laws

Equation (5) matches the Monte Carlo gain loss to within 0.007 dB at B=1 and within 0.0003 dB for B≥2, and DR lies 0.01–0.09 dB above it (Figure 6a). QACD lies above both (0.36 versus 0.22 dB at B=3) because it deliberately trades a little gain for lower sidelobes, while RPO+QACD stays at 0.24 dB. Equation (8) matches the Monte Carlo pointing error to within 4% for B≥2 and underestimates it by 17% at B=1, where the linearisation of (7) breaks down (Figure 6b). DR is 1.5–1.7 times worse than (8) for B≥2, the correlated-error effect described in the error-law analysis. The QACD pointing error stays between 3×10^{-4} and 5×10^{-4} in u, about 0.3 u_{tol}, because with ρ=1 the optimiser accepts a small error whenever this buys PSL.

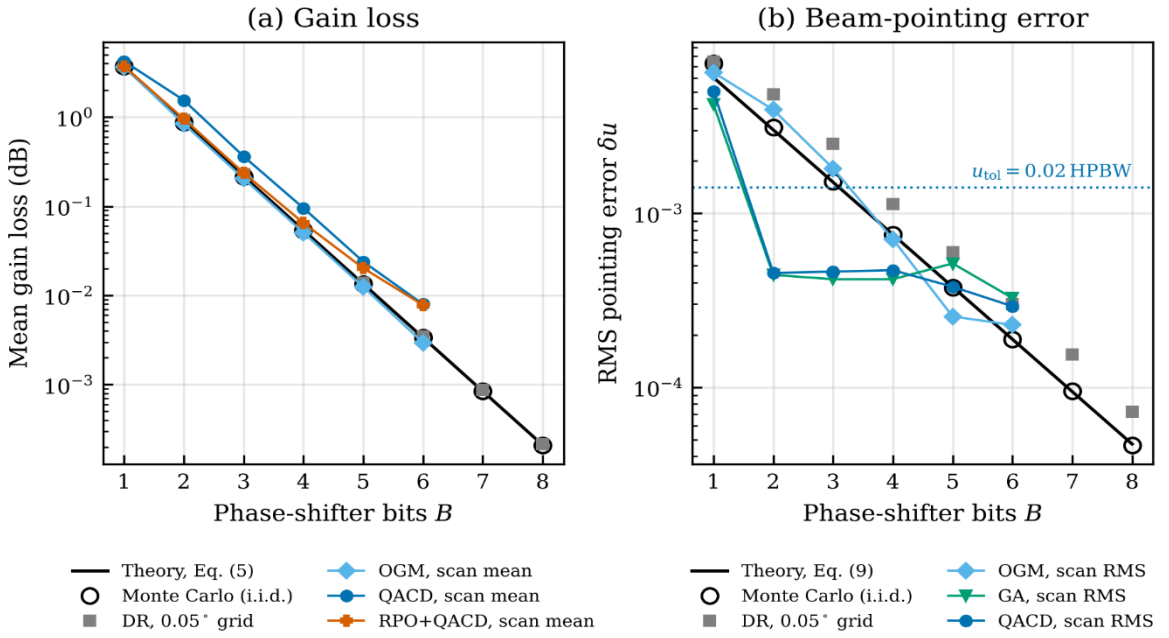

FIGURE 6: Validation of the Error Laws Versus B. (a) Mean Gain Loss: Theory (5), Monte Carlo with Independent Errors (20 000 Trials per Point), DR on a 0.05° Scan Grid, and Scan Means of OGM, QACD, and RPO+QACD. (b) RMS Pointing Error in u: Theory (8), Monte Carlo, DR, and Scan RMS of OGM, GA, and QACD; the Dotted Line Is u_{tol}.

Pointing Accuracy and Beamwidth Across the Scan

The pointing errors of DR and OGM are concentrated at a few angles and reach 0.81° for DR at θ_0=49° (Figure 7a), where the error sequence has a long period and a large linear component. QACD keeps all 61 commanded angles within ±2% of the HPBW (Figure 7b), GA 60 of them, and DR and OGM only 80% and 84%. Quantization changes the beamwidth very little, so it is reported in Table 3 rather than in a separate figure: QACD broadens the beam by 0.060° on average at B=3 (at most 0.22°).

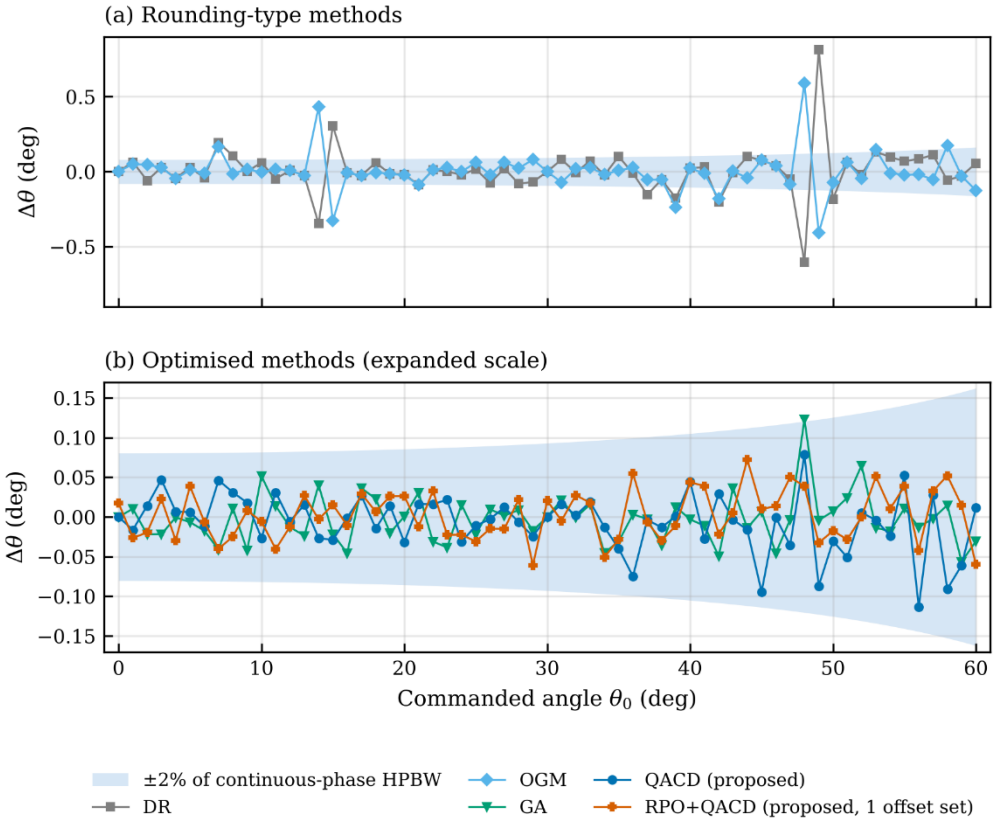

FIGURE 7: Beam-Pointing Error Versus Commanded Angle at B=3. (a) DR and OGM. (b) GA, QACD and RPO+QACD on an Expanded Scale. The Shaded Band Is ±2% of the Continuous-Phase HPBW.

Planar Array

The 16 × 16 array shows the largest benefit, because 256 elements give QACD many more degrees of freedom (Figure 8). At B=3, the PSL improves from −16.9 dB (DR) to −27.8 dB, within 2.2 dB of the −30.0 dB continuous design, while the gain loss rises from 0.22 to 0.48 dB. At B=2, the PSL improves from −14.0 to −23.1 dB, with the gain loss rising from 0.91 to 1.35 dB. QACD took 3.6 s (B=3) and 2.0 s (B=2) per beam, acceptable for an offline beam table.

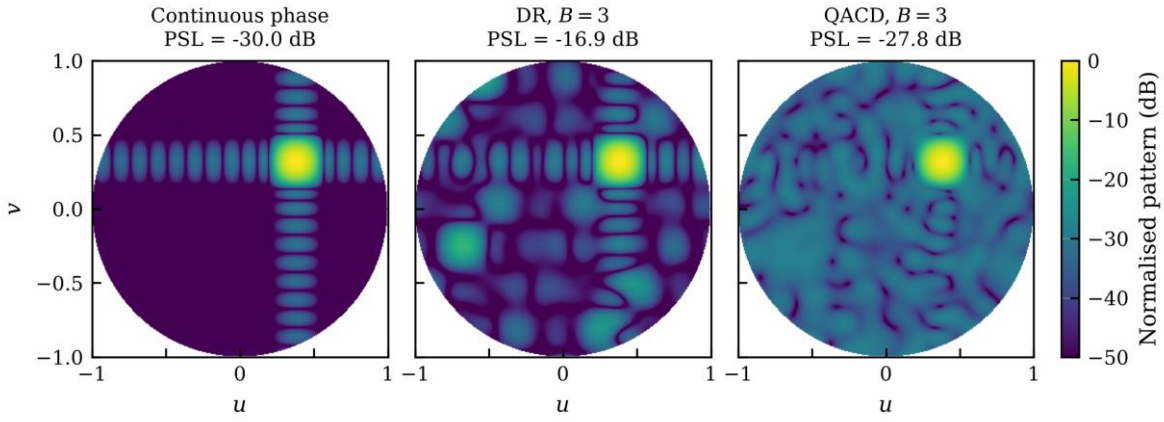

FIGURE 8: Normalised $u - v$ Patterns of the 16 × 16 Array Steered to (θ_0, φ_0)=(30°, 40°) at B=3: Continuous Phase, DR and QACD.

Convergence and Complexity

At θ_0=23°, QACD improves on its OGM starting point within tens of evaluations and stops after 640, 1280 and 3072 evaluations for B=2, 3 and 4 (Figure 9). GA, after 3840 evaluations, ends 0.9 dB lower in cost at B=2 and 0.2 dB lower at B=4, but 0.1 dB higher at B=3.

Without continuation, QACD stops earlier; at this angle it ends 1.4 and 0.2 dB higher for B=3 and 4 but 1.1 dB lower for B=2. A single angle therefore does not settle the value of continuation, but the scan averages do: continuation improves the mean PSL at every B≥2 and the worst case even more (−19.3 versus −17.5 dB at B=3; Table 3).

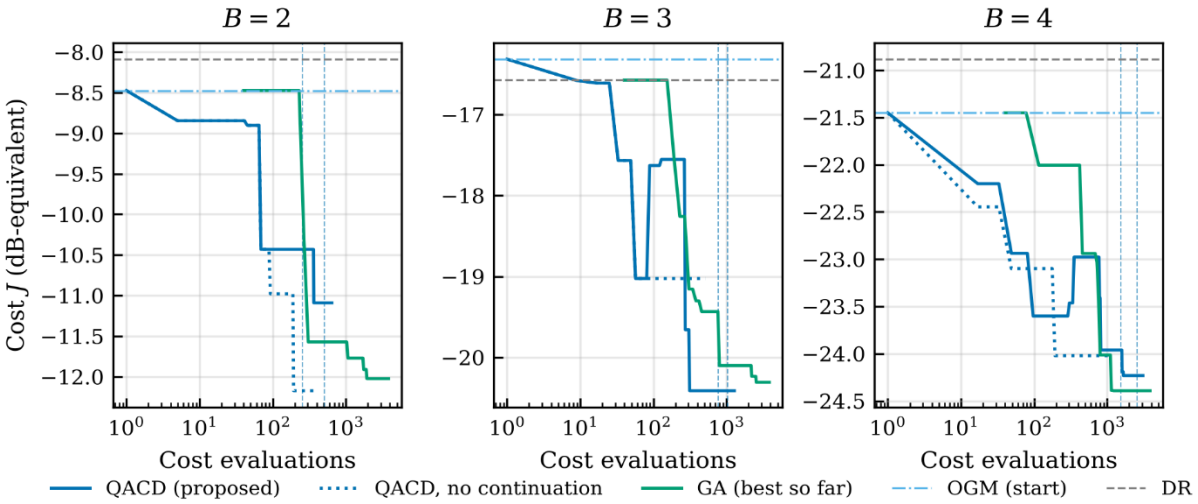

FIGURE 9: True Cost J Versus Number of Cost Evaluations at θ_0=23° for B=2, 3, and 4. Vertical Dashed Lines Mark the Continuation Stage Switches; GA Shows the Best Cost Found So Far.

Sensitivity to the Cost Weights

The pointing weight ρ is the real trade-off knob (Figure 10a). Without a pointing term (ρ=0), QACD reaches −22.6 dB mean PSL but leaves a 0.35° RMS pointing error, about twice that of DR; ρ=1 gives −21.7 dB and 0.039°, and ρ=100 gives 0.005° but only −18.1 dB. By (7), the pointing error lives in the linear component of the phase error, which also shapes the near-in sidelobes, so constraining it removes freedom the optimiser would otherwise use against the sidelobes. The gain weight κ is almost free (Figure 10b): raising it from 0 to 16 lowers the mean gain loss from 0.34 to 0.28 dB while the mean PSL stays within ±0.1 dB of −21.7 dB. Gain loss depends on the mean of $\cos\varepsilon_n$, whereas PSL depends on how the error energy is distributed in angle, so the two barely compete. The pre-set κ=1 is therefore not the best choice and κ≈4–16 is recommended; to avoid tuning on the reported data, all other results keep κ=1 and Table 3 adds post-hoc rows with κ=4. The values in this subsection come from the 3° scan step used for the sensitivity studies (Table 2), so they differ slightly from the 1° results of Table 3 (for example −21.7 versus −21.0 dB mean PSL at ρ = κ = 1). Because κ = 4 changes the mean PSL by only 0.03 dB in Table 3, retaining the pre-set κ = 1 for the headline results does not alter any conclusion.

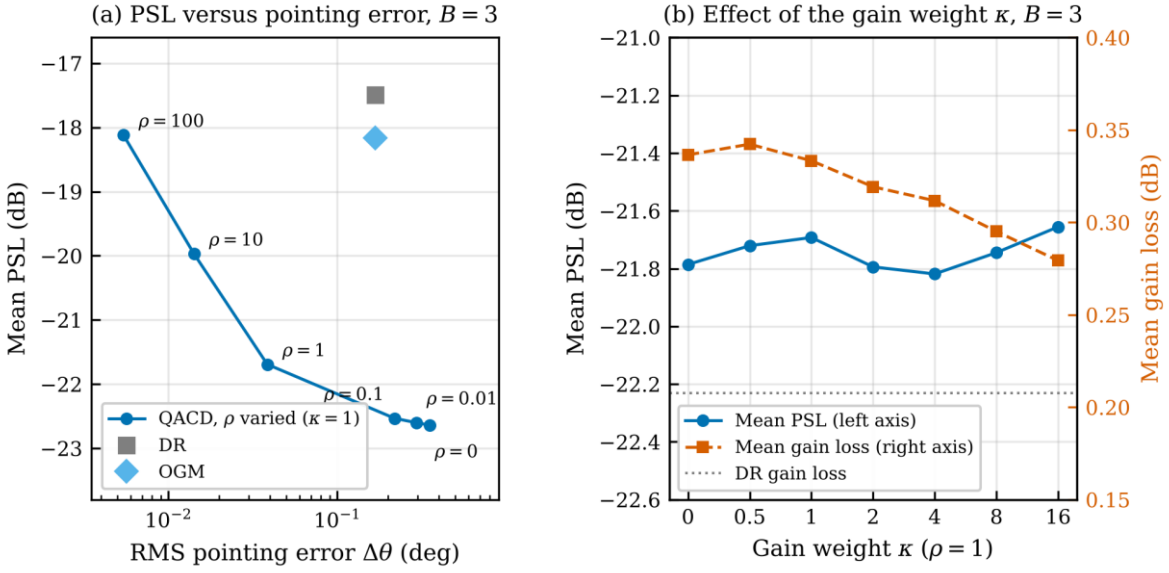

FIGURE 10: QACD Trade-Offs at B=3 (θ_0=0°–60°, 3° Step). (a) Mean PSL Versus RMS Pointing Error as ρ Varies (κ=1), with DR and OGM for Reference. (b) Mean PSL (Left Axis) and Mean Gain Loss (Right Axis) Versus κ (ρ=1); the Dotted Line Is the Gain Loss of DR.

Robustness to Hardware Error and Mutual Coupling

The benefit of any deterministic state selection erodes as the random hardware error approaches the quantization error itself, $\Delta/\sqrt{12}$=13° at B=3 (Figure 11a). The advantage of QACD over DR falls from 4.2 dB without hardware error to 2.4 dB at 8° RMS and 0.8 dB at 15°, so the phase-shifter calibration must be several times finer than the LSB for state optimisation to pay off. Under the Toeplitz coupling model, the continuous-phase PSL hardly moves (−30.2 to −29.6 dB at α=0.3), but the optimised states are sensitive to it (Figure 11b): optimised without knowledge of C, QACD loses 2.4 dB of PSL at α=0.3, whereas optimised with C it keeps −21.0 dB.

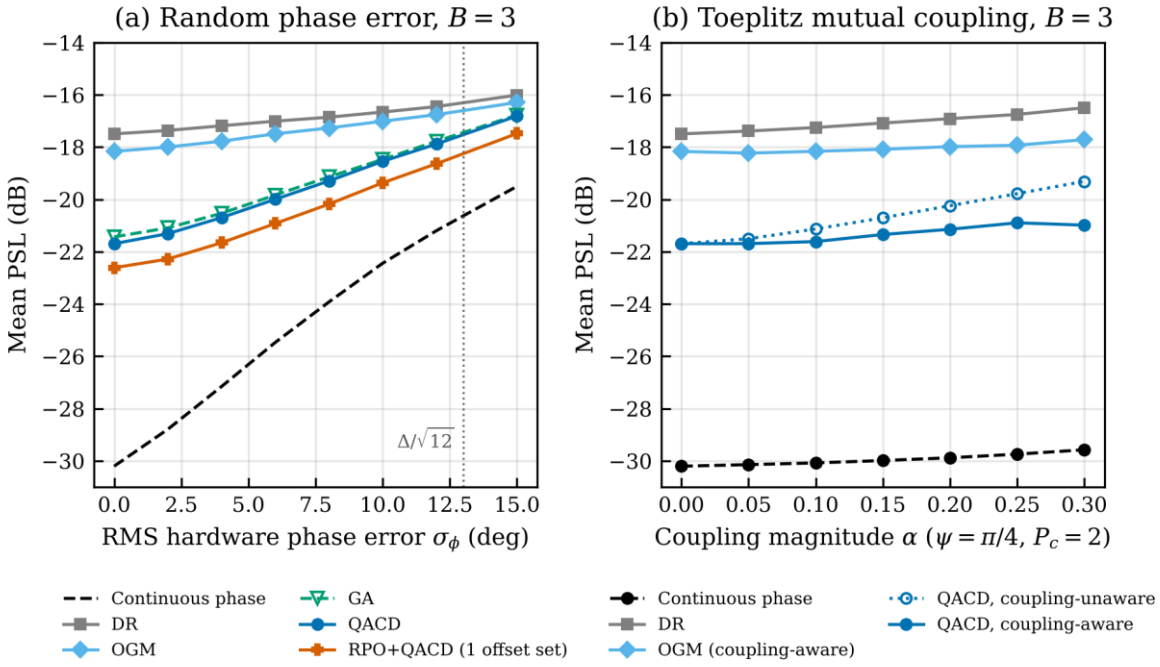

FIGURE 11: (a) Mean PSL Versus RMS Random Phase Error Added to the Selected States (B=3, 200 Trials per Point). (b) Mean PSL Versus Coupling Magnitude α in the Toeplitz Model (B=3). Both Panels: θ_0=0°–60°, 3° Step.

Effect of Array Size

DR barely improves with array size, because its quantization lobes are set by the number of bits rather than by N. QACD does improve (Figure 12a): at B=3 its advantage over DR grows from 2.3 dB at N=16 to 4.0, 6.0 and 8.8 dB at N=32, 64 and 128, because more elements give more independent ways to cancel a lobe. The run time grows from 12 to 217 ms per beam at B=3, roughly as $N^{1.4}$ (Figure 12b).

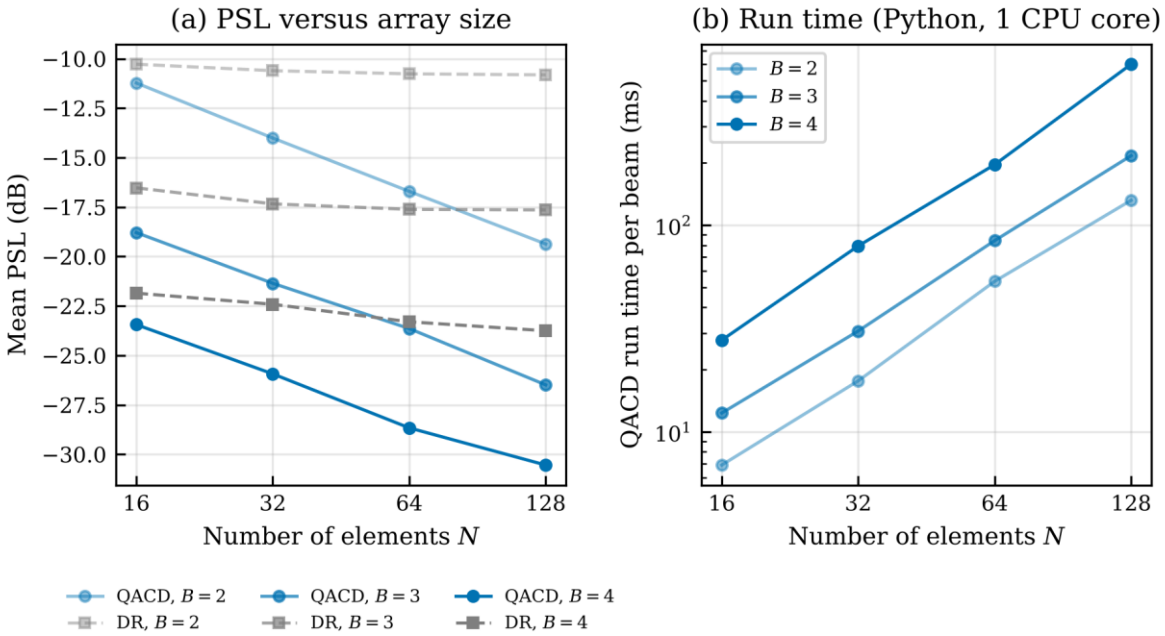

FIGURE 12: (a) Mean PSL Versus N (θ_0=0°–60°, 4° Step) for B=2, 3, and 4; QACD Solid, DR Dashed, Darker Lines for More Bits. (b) QACD Run Time per Beam.

Comparative Analysis

No method wins on every metric (Tables 3 and 4), and three operating regimes emerge. The gain-first methods (DR, OGM) give the best gain and the narrowest beam, but quantization lobes 4–5 dB above those of the optimised methods and pointing errors of up to 0.8°; OGM improves the mean gain over DR by only 0.02 dB, so exact gain optimality matters little for a steered array whose ideal phases are already nearly aligned. The peak-first methods (QACD, GA) lower the mean PSL by about 4 dB and the RMS pointing error four- to five-fold, at a cost of 0.14–0.15 dB of gain and a 1.4–1.5 dB higher ISL. The ISL increase follows from energy conservation: with $\sum A_n^2$ fixed, power removed from the main beam must appear in the sidelobes, and minimising the largest lobe spreads it evenly. Hardware plus optimisation (RPO+QACD) gives the best PSL (−22.4 dB mean, −20.9 dB worst) with DR-level gain loss and the lowest ISL, because the fixed pre-phases already supply the decorrelation that QACD otherwise creates from the states. This agrees with the conclusion that pre-phasing and state selection are complementary [11]. Direct numerical comparison with published methods is not meaningful, because their arrays, tapers, scan sets and metrics differ.

TABLE 3: Beam Quality at B=3, N=32, Taylor −30 dB, θ_0=0°–60° (1° Step).

Method	Mean PSL (dB)	Worst PSL (dB)	ISL (dB)	Mean GL (dB)	RMS $\Delta\theta$ (°)	Max $\lvert\Delta\theta\rvert$ (°)	Mean $\lvert\Delta$HPBW$\rvert$ (°)
Continuous phase (ref.)	−30.20	−30.20	−23.26	0	0	0	0
RPO+QACD (proposed)	**−22.38**	**−20.85**	**−13.28**	0.236	0.032	0.081	0.048
QACD (proposed)	−21.02	−19.28	−11.49	0.362	0.037	0.114	0.060
GA	−20.81	−18.65	−11.41	0.369	0.030	0.123	0.056
QACD-nc (ablation)	−20.38	−17.51	−11.14	0.394	**0.028**	**0.066**	0.069
RPO	−20.46	−17.52	−12.98	0.218	0.113	0.323	0.014
OGM	−17.09	−13.95	−13.15	**0.204**	0.133	0.590	**0.010**
DR	−16.70	−13.95	−12.90	0.221	0.161	0.815	**0.010**
QACD, κ = 4 (post hoc)	−21.05	−19.44	−11.59	0.337	0.033	0.096	0.042
RPO+QACD, κ = 4 (post hoc)	−22.42	−21.05	−13.38	0.220	0.031	0.072	0.036

Best quantized value per column (κ = 1 rows) in bold. ISL is averaged over the scan in linear power. RPO rows average 10 offset realisations (RPO+QACD: 3). Post-hoc rows use κ = 4.

TABLE 4: Cost of Each Method at B=3, N=32.

Method	Cost evaluations	Time (ms)	Extra hardware
RPO+QACD (proposed)	1609	28	Pre-phases
QACD (proposed)	1859	33	None
GA	3840	134	None
QACD-nc (ablation)	919	10	None
RPO	1	<1	Pre-phases
OGM	256	0.5	None
DR	1	<1	None

Mean per beam over 61 scan angles; Python wall-clock time on one core. QACD figures are for the main seed; Table 5 gives the mean over ten seeds (1810 evaluations, 32 ms).

Because both QACD and GA depend on a random seed, their comparison was repeated with ten seeds (Table 5). The two methods reach the same mean PSL: the 0.09 dB difference is about one standard deviation. QACD has a 0.5 dB better worst case on average and a smaller spread, needs 2.1 times fewer cost evaluations, and runs about four times faster in this implementation. GA remains a strong offline baseline, being better at B=2 and at many individual angles. For B≥5, the L-fold enumeration makes QACD costlier than the fixed-budget GA (14 710 versus 3840 evaluations at B=6), although QACD is still 0.4 dB better in mean PSL there.

TABLE 5: QACD Versus GA over Ten Random Seeds (B=3, θ_0=0°–60°, 1° Step).

Method	Mean PSL (dB)	Worst PSL (dB)	Mean GL (dB)	Cost evaluations	Time (ms)
QACD	−20.98 ± 0.06 [−21.06, −20.88]	−19.23 ± 0.28 [−19.55, −18.77]	0.365	1810	32
GA	−20.89 ± 0.08 [−21.02, −20.75]	−18.70 ± 0.44 [−19.27, −17.91]	0.365	3840	137

Mean ± sample standard deviation over seeds, with the range in brackets.

Effect of the Amplitude Taper

With a −35 dB Taylor taper the main result is repeated (Table 6): QACD improves the mean PSL over DR by 4.5 dB and the worst case by 4.2 dB at a similar gain cost. With a uniform taper the −13.2 dB first sidelobes already dominate, so there is no quantization lobe to remove; QACD instead pushes the PSL to −14.7 dB by using the phases as a crude phase-only taper, at a cost of 0.85 dB of gain and a 2.2 dB higher ISL. QACD is therefore most useful with a taper whose design sidelobes lie well below the quantization lobes.

TABLE 6: Effect of the Amplitude Taper (N=32, B=3, θ_0=0°–60°, 1° Step).

Method	Mean PSL (dB)	Worst PSL (dB)	ISL (dB)	Mean GL (dB)	RMS $\Delta\theta$ (°)
Taylor −35 dB					
Continuous	−35.18	−35.18	−27.60	0	0
DR	−17.07	−13.99	−13.19	0.222	0.173
OGM	−17.74	−13.99	−13.43	0.204	0.134
QACD	−21.57	−18.14	−11.82	0.365	0.039
Uniform					
Continuous	−13.23	−13.23	−10.14	0	0
DR	−12.25	−9.36	−8.57	0.216	0.138
OGM	−12.42	−9.36	−8.63	0.202	0.124
QACD	−14.65	−12.99	−6.33	0.854	0.019

Limitations

The results describe the array-processing effect of phase quantization under an idealised model and are not predictions for a specific antenna. Element patterns, embedded-element effects, edge effects and scan blindness are not modelled. The phase states are ideal, without state-dependent insertion loss or phase error, so a calibrated state table should replace them before QACD is applied; random errors comparable to the LSB remove most of the benefit (Figure 11a). The coupling model is an analytical stress test that assumes **C** is known exactly. Beam squint and true-time-delay steering are not considered. QACD stops at a local minimum whose quality depends on the element visiting order. Most comparisons use one array and one u_{tol}, and run times come from unoptimised Python. Minimising PSL raises the ISL by about 1.4 dB without pre-phases, so applications limited by total interference may prefer DR or OGM. The joint cost is a scalarisation, so no Pareto optimality is claimed, and the variability caused by the element visiting order was quantified only at B = 3 (Table 5). Finally, no hardware was built or measured; validation on a low-bit array is the most important next step.

CONCLUSIONS

Choosing the states of low-bit phase shifters with a joint cost, rather than rounding each phase independently, lowers the worst-case quantization lobe and the RMS pointing error substantially for a fraction of a decibel of gain. For a 32-element Taylor array with 3-bit phase shifters, QACD improved the mean PSL by 4.3 dB and the worst-case PSL by 5.3 dB, reduced the RMS pointing error 4.4-fold and cost 0.14 dB of additional gain loss relative to direct rounding; with fixed random pre-phases the mean PSL improved by 5.7 dB at almost no gain cost, and on a 16 × 16 planar array the PSL improved by 10.9 dB. QACD reached the same mean PSL as a genetic algorithm with half the cost evaluations.

The first-order laws explain these results: gain loss is governed by the mean of $\cos\varepsilon_n$ and falls fourfold per bit, whereas the pointing error is governed by the linear component of the phase error and halves per bit. Three practical rules follow. A 1-bit array without pre-phases cannot avoid a 0 dB mirror lobe; for 2–4 bits and a tapered aperture, state optimisation is worthwhile and its benefit grows with array size; and it pays off only if the calibration error is well below the quantization error. Future work should include measured state-dependent phase-shifter errors, embedded element patterns, wideband operation and hardware validation.

ACKNOWLEDGMENT

The authors thank the School of Electronics and Information Engineering, Nanjing University of Information Science and Technology (NUIST), for its support. This research received no external funding.

REFERENCES

[1] Mailloux, R. J. (2005). Phased array antenna handbook (2nd ed.). Norwood, USA: Artech House.

[2] Wang, Z., Li, M., Liu, Q., & Swindlehurst, A. L. (2018). Hybrid precoder and combiner design with low-resolution phase shifters in mmWave MIMO systems. IEEE Journal of Selected Topics in Signal Processing, 12(2), 256–269. https://doi.org/10.1109/JSTSP.2018.2819129

[3] Lyu, S., Wang, Z., Gao, Z., He, H., & Hanzo, L. (2021). Lattice-based mmWave hybrid beamforming. IEEE Transactions on Communications, 69(7), 4907–4920. https://doi.org/10.1109/TCOMM.2021.3075248

[4] Shao, F., Wang, H., Huang, C., Li, G., & Gong, W. (2025). A novel phased array architecture with mixed-resolution phase shifters.

IEICE Transactions on Communications, E108-B(6), 733–740. https://doi.org/10.23919/transcom.2024EBP3132

[5] Kashyap, B. G., Theofanopoulos, P. C., Cui, Y., & Trichopoulos, G. C. (2020). Mitigating quantization lobes in mmWave low-bit reconfigurable reflective surfaces. IEEE Open Journal of Antennas and Propagation, 1, 604–614. https://doi.org/10.1109/OJAP.2020.3034049

[6] Vabichevich, D., Belov, A., & Sayanskiy, A. (2023). Suppression of quantization lobes in 1-bit reconfigurable intelligent surfaces. IEEE Antennas and Wireless Propagation Letters, 22(12), 2808–2811. https://doi.org/10.1109/LAWP.2023.3299198

[7] Shekhawat, A. S., Kashyap, B. G., Raldiris Torres, R. W., Shan, F., & Trichopoulos, G. C. (2025). A millimeter-wave single-bit reconfigurable intelligent surface with high-resolution beam-steering and suppressed quantization lobe. IEEE Open Journal of Antennas and Propagation, 6(1), 311–325. https://doi.org/10.1109/OJAP.2024.3506453

[8] Keshmiri, S., Jayasuriya, S., & Imani, M. F. (2026). Novel 1-bit hybrid reconfigurable intelligent surface. Scientific Reports, 16, 25976. https://doi.org/10.1038/s41598-026-55424-w

[9] Smith, M. S., & Guo, Y. C. (1983). A comparison of methods for randomizing phase quantization errors in phased arrays. IEEE Transactions on Antennas and Propagation, 31(6), 821–828. https://doi.org/10.1109/TAP.1983.1143152

[10] Jiang, W., Guo, Y., Liu, T., Shen, W., & Cao, W. (2003). Comparison of random phasing methods for reducing beam pointing errors in phased array. IEEE Transactions on Antennas and Propagation, 51(4), 782–787. https://doi.org/10.1109/TAP.2003.811105

[11] Cheng, M., Wu, Q., Yu, C., Wang, H., & Hong, W. (2023). A prephased electronically steered phased array that uses very-low-resolution phase shifters and a hybrid phasing method. IEEE Transactions on Antennas and Propagation, 71(9), 7310–7322. https://doi.org/10.1109/TAP.2023.3291778

[12] Krieger, J. D., Yeang, C. P., & Wornell, G. W. (2013). Dense delta-sigma phased arrays. IEEE Transactions on Antennas and Propagation, 61(4), 1825–1837. https://doi.org/10.1109/TAP.2013.2241719

[13] Zhang, Y., Zhao, G., Sun, H., He, M., & Chen, Q. (2016). Linear programming phase feeding method for phased-array scanning. IEICE Transactions on Electronics, E99-C(7), 892–894. https://doi.org/10.1587/transele.E99.C.892

[14] Zhang, Y., Sun, H., & Liu, Q. H. (2019). Integer linear programming roundoff method for arbitrary planar phased-array scanning. IEEE Transactions on Antennas and Propagation, 67(5), 3040–3047. https://doi.org/10.1109/TAP.2019.2896734

[15] Thompson, D., Yeary, M., Fulton, C., & McGuire, B. (2015). Optimized beam steering approach for improved sidelobes in phased array radars using a minimal number of control bits. IEEE Transactions on Antennas and Propagation, 63(1), 106–112. https://doi.org/10.1109/TAP.2014.2366198

[16] Zhang, Y., Shen, K., Ren, S., Li, X., Chen, X., & Luo, Z. Q. (2022). Configuring intelligent reflecting surface with performance guarantees: optimal beamforming. IEEE Journal of Selected Topics in Signal Processing, 16(5), 967–979. https://doi.org/10.1109/JSTSP.2022.3176479

[17] Ren, S., Shen, K., Li, X., Chen, X., & Luo, Z. Q. (2023). A linear time algorithm for the optimal discrete IRS beamforming. IEEE Wireless Communications Letters, 12(3), 496–500. https://doi.org/10.1109/LWC.2022.3232146

[18] Saleem, R. A. B., Shah, A. A., Munsif, H., Najam, A. I., Khattak, S., & Irfanullah. (2024). Pattern compensation of a planar phased array with centre elements phase malfunctioning using a genetic algorithm. Progress In Electromagnetics Research Letters, 122, 21–28. https://doi.org/10.2528/PIERL24041404

[19] Ge, Q., Zhang, Y., & Wang, Y. (2020). A low complexity algorithm for direction of arrival estimation with direction-dependent mutual coupling. IEEE Communications Letters, 24(1), 90–94. https://doi.org/10.1109/LCOMM.2019.2950029

[20] Harris, C. R., Millman, K. J., van der Walt, S. J., Gommers, R., Virtanen, P., Cournapeau, D., et al. (2020). Array programming with NumPy. Nature, 585(7825), 357–362. https://doi.org/10.1038/s41586-020-2649-2

[21] Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., et al. (2020). SciPy 1.0: fundamental algorithms for scientific computing in Python. Nature Methods, 17, 261–272. https://doi.org/10.1038/s41592-019-0686-2