

Pair-Structured Polynomial Rooting in Signal Processing: Polytopic Cascade Design, A Spread Identity for Root-MUSIC, and What It Says About Breakdown Detection

Sawon Pratiher

sawon1234@gmail.com

Home, Andul, Howrah. <https://orcid.org/0000-0001-8471-8148>

Research Article

Keywords: Root-MUSIC, direction-of-arrival estimation, polynomial rooting, second-order sections, IIR filter design, Schur stability, threshold effect, subspace methods, performance breakdown

Posted Date: October 27th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11497536/v1>

License:  This work is licensed under a Creative Commons Attribution 4.0 International License.
[Read Full License](#)

Additional Declarations: The authors declare no competing interests.

Pair-Structured Polynomial Rooting in Signal Processing: Polytopic Cascade Design, a Spread Identity for Root-MUSIC, and What It Says About Breakdown Detection

Sawon Pratiher

Abstract—Two common polynomial-rooting tasks in signal processing, cascade IIR design and subspace direction-of-arrival estimation, carry symmetries that standard rooting discards. We study both through a *pair-structured* parameterization in which the polynomial is a product of second-order sections, each described by a *centre* and a *spread*. For cascade design the Schur region of a section is the triangle $|c| < 1$, $|2A| < 1 + c$, so the stable region of a whole cascade is a polytope; magnitude fitting subject to guaranteed stability then becomes a smooth problem over a convex set, and we observe roughly an order of magnitude of error reduction per added section. For root-MUSIC the relevant symmetry is conjugate-reciprocal, roots occurring in pairs $(z, 1/z^*)$ that *coalesce on the unit circle* in the noiseless limit because the MUSIC null spectrum is non-negative with double zeros at the true angles. This degenerate-pair structure is the central object of the paper. We prove an exact first-order identity, $v_k^2 = 2f_{\min}/f''$, relating the pair spread $v_k = \log |z_k|$ to the residual null-spectrum depth and its curvature, and we verify the resulting scaling law $v \sim N^{-1/2}$ to three decimal places across four SNRs. We report three findings, two of them negative and stated plainly because each contradicts a natural expectation. First, the pair-symmetrized angle estimator is *algebraically identical* to textbook root-MUSIC: symmetry exploitation yields no variance reduction at all, and our Monte Carlo confirms bit-identical RMSE at every SNR. Second, the pair spread is a strong breakdown detector, reaching area under the ROC curve of 0.974 at -6 dB against 0.842 for the eigenvalue gap and 0.575 for an AIC statistic, at literally zero additional cost. Third, and against our own framing, the raw null depth $f_{\min}$ *outperforms* the spread uniformly, reaching AUC 1.000 at -6 dB; our identity explains why, since dividing by a trial-varying curvature injects variance. The practical recommendation that follows is therefore not the one we set out to make.

Index Terms—Root-MUSIC, direction-of-arrival estimation, polynomial rooting, second-order sections, IIR filter design, Schur stability, threshold effect, subspace methods, performance breakdown.

I. INTRODUCTION

POLYNOMIAL rooting appears throughout signal processing. In cascade IIR realization a transfer function is implemented as a product of biquads, both for coefficient-sensitivity reasons and because the second-order section is the natural fixed-point building block [1], [2]. In array processing, root-MUSIC [4], [5] replaces the spectral search of MUSIC [3] by rooting a polynomial built from the noise subspace, giving lower complexity and better accuracy for uniform linear arrays.

In both settings the polynomial carries a symmetry. A real filter polynomial factors over $\mathbb{R}$ into second-order sections.

The root-MUSIC polynomial is conjugate-reciprocal, its roots occurring in pairs $(z, 1/z^*)$. We ask what is gained by a parameterization respecting this structure, in which each pair is described by a *centre* and a *spread* rather than by two independent roots. The underlying algebra, in which the spread sign classifies the pair and the discriminant factorizes accordingly, is developed in [11].

A. Contributions

- 1) **Polytopic cascade stability.** In pair coordinates (A, c) for a section $z^2 - 2Az + c$, the Schur region is the triangle with vertices $(1, 1)$, $(-1, 1)$, $(0, -1)$ (Theorem 1), verified on 2×10^5 random sections with no mismatches and confirmed convex on 5×10^4 tests. The cascade region is the Cartesian power. A design study (Section II-A) shows near-geometric error decay with the number of sections and all sections strictly feasible throughout.
- 2) **The degenerate-pair structure of root-MUSIC.** Because the MUSIC null spectrum is a non-negative trigonometric polynomial, its zeros on the unit circle are *double*, so each signal root pair $(z_k, 1/z_k^*)$ coalesces in the noiseless limit and splits radially under noise (Proposition 2). This identifies the pair spread as the splitting of a double root, and places root-MUSIC in the degenerate case $\Lambda = 0$ of the pair-partition classification.
- 3) **An exact spread identity.** We prove $v_k^2 = 2f_{\min}/f''(\hat{\omega}_k) + O(v^4)$ (Theorem 2), verified with median relative deviation 0.19 to 0.28 across six SNR and sample-size settings, together with the consequent scaling $v \sim N^{-1/2}$, verified as $N^{-0.505}$, $N^{-0.492}$, $N^{-0.494}$, $N^{-0.487}$ at four SNRs.
- 4) **Two negative results, stated plainly.** Pair symmetrization gives exactly zero accuracy gain (Proposition 3), and the raw null depth $f_{\min}$ beats the spread as a breakdown detector at every SNR tested (Section V-B). We report both because each contradicts a natural expectation and because our own identity predicts the second.

We do not claim improved DOA accuracy. Root-MUSIC is already asymptotically efficient for uncorrelated sources on a uniform linear array [5], [6], and our Monte Carlo confirms it attains the stochastic Cramér-Rao bound above threshold.

II. PAIR COORDINATES AND CASCADE STABILITY

Definition 1 (Pair coordinates). A monic second-order section is written

$$Q(z) = z^2 - 2Az + c = (z - A)^2 - \Lambda, \quad \Lambda = A^2 - c, \quad (1)$$

Algorithm 1 Cascade design in pair coordinates

Require: target $|H_d|$ on a grid, section count K , starts R
Ensure: gain g , numerators $b^{(k)}$, sections $(A_k, c_k) \in \mathcal{T}$ for
 $t = 1, \dots, R$ do
 l:
 end
 draw (A_k, c_k) inside $\mathcal{T}$ and $b^{(k)}$ at random
 2: minimize $\sum_i (|H(\omega_i)| - |H_d(\omega_i)|)^2 + \rho \sum_k \pi_{\mathcal{T}}(A_k, c_k)$
 3: where $\pi_{\mathcal{T}}(A, c) = (|c| - 1)_+^2 + (|2A| - (1 + c))_+^2$
 4:
 5: **return** best feasible design = 0

with A the *pair centre*, c the *pair product* and Λ the *pair spread*. The roots are $A \pm \sqrt{\Lambda}$: a real pair when $\Lambda > 0$, a repeated root when $\Lambda = 0$, a conjugate pair when $\Lambda < 0$.

Theorem 1 (Polytopic cascade stability). *Section (1) with real A, c has both roots strictly inside the unit circle iff*

$$|c| < 1 \quad \text{and} \quad |2A| < 1 + c, \quad (2)$$

the open triangle $\mathcal{T}$ with vertices $(A, c) = (1, 1), (-1, 1), (0, -1)$. For a cascade $H(z) = g \prod_{k=1}^K B_k(z)/Q_k(z)$ the set of stable denominator parameters is the polytope $\mathcal{T}^K$.

Proof: Writing $Q(z) = z^2 + a_1 z + a_0$ with $a_1 = -2A$, $a_0 = c$, the Jury test for degree two gives $|a_0| < 1$ and $|a_1| < 1 + a_0$, which is (2). The three constraints $c < 1$, $2A < 1 + c$, $-2A < 1 + c$ are affine, so $\mathcal{T}$ is a polyhedron; solving the three pairs of equalities gives the stated vertices. A product is Schur iff each factor is, and no constraint couples sections. ■

We verified (2) against direct root computation on 2×10^5 sections with (A, c) uniform on $[-1.5, 1.5]^2$, with zero mismatches, and verified convexity on 5×10^4 random convex combinations of admissible points, with zero violations.

Remark 1. The triangle is classical. What Theorem 1 adds for design is the product structure: the constraint is separable across sections and convex in each, so cascade design is a smooth fit over a polytope with no stability projection or post-hoc pole reflection. The sign of Λ_k additionally reports, for free, whether section k realizes a real pole pair or a resonant pair, which is what a fixed-point implementer needs for wordlength allocation.

The penalty $\pi_{\mathcal{T}}$ is a squared distance to a triangle, hence convex in (A_k, c_k) and separable across k ; contrast a penalty on the largest pole modulus in coefficient space, which requires root-finding to evaluate and is neither convex nor separable.

A. Cascade design study

We fit the magnitude target $|H_d(\omega)| = (1 + (\omega/\omega_c)^{2L})^{-1/2}$ with $\omega_c = 0.3\pi$ on a grid of 400 frequencies using K cascaded biquads parameterized directly by (A_k, c_k) and numerator coefficients, with a penalty enforcing (2), from 25 random starts. Table I reports the resulting fit error as K grows with the target order $L = 2K$.

Table I: Cascade magnitude fit: error versus number of sections. All designed sections lie strictly inside the Schur triangle at every K .

K	target order	max error	rms error	all in $\mathcal{T}$
1	2	1.064×10^{-2}	2.408×10^{-3}	yes
2	4	9.915×10^{-4}	1.646×10^{-4}	yes
3	6	8.811×10^{-5}	1.212×10^{-5}	yes
4	8	2.528×10^{-6}	3.625×10^{-7}	yes

Table II: Three-section lowpass design, 40 random starts per parameterization. “Stable” counts designs whose poles all lie inside the unit circle.

parameterization	best max err	median max err	stable
pair coordinates	2.52×10^{-5}	8.81×10^{-5}	40/40
direct coefficients	8.81×10^{-5}	1.196	1/40

The maximum error falls by factors of 10.7, 11.3 and 34.9 as sections are added, roughly geometric decay. The designed three-section cascade has (A_k, c_k) equal to $(0.4093, 0.2637)$, $(0.4871, 0.2353)$ and $(0.4807, 0.6139)$, with spreads -0.0962 , $+0.0020$ and -0.3828 : the spread signs correctly report that sections 1 and 3 are resonant while section 2 sits essentially at the critically damped boundary.

B. Why the parameterization matters

It is reasonable to ask whether the pair coordinates buy anything that a direct fit in denominator-coefficient space would not. They buy a great deal. We repeated the three-section design from 40 random starts in two parameterizations: the pair coordinates (A_k, c_k) with the separable penalty enforcing (2), and a direct sixth-order rational fit in the coefficients $(b_0, \dots, b_6, a_1, \dots, a_6)$ with a penalty on the largest pole modulus. Table II reports the outcome.

The gap is not marginal. In pair coordinates every one of the 40 runs returned a stable design, and the median fit error is within a factor of four of the best achieved. In coefficient space only one run in forty returned a stable filter, and the median error of 1.196 indicates that most runs failed outright: the optimizer wanders into the unstable region, where the penalty is a non-convex function of the coefficients computed through a root-finding step, and does not recover. The separability and convexity of (2) is what removes this failure mode, since each section can be kept feasible independently and the feasible set has no re-entrant geometry for the optimizer to become trapped against.

Remark 2. This comparison isolates the parameterization, not the algorithm: both runs used the same Levenberg-Marquardt solver, the same frequency grid, the same target and the same number of starts. The only difference is the coordinates in which stability is expressed.

III. THE DEGENERATE PAIR STRUCTURE OF ROOT-MUSIC

A. Signal model

Consider a uniform linear array of M sensors at half-wavelength spacing with $d < M$ narrowband sources at

Algorithm 2 Pair-structured root-MUSIC with breakdown flag

Require: snapshots $\mathbf{x}(1:N)$, source count d , threshold τ
Ensure: angles $\hat{\theta}_{1:d}$, spreads $\nu_{1:d}$, flag

```

1:  $\hat{\mathbf{R}} \leftarrow N^{-1} \sum_t \mathbf{x}(t) \mathbf{x}^H(t)$   $\triangleright O(NM^2)$ 
2: eigendecompose  $\hat{\mathbf{R}}$ ; take  $\hat{\mathbf{E}}_n$   $\triangleright O(M^3)$ 
3:  $\mathbf{C} \leftarrow \hat{\mathbf{E}}_n \hat{\mathbf{E}}_n^H$ ;  $c_l \leftarrow$  diagonal sums  $\triangleright O(M^2)$ 
4: root  $z^{M-1} p(z)$   $\triangleright O(M^3)$ 
5: group roots into pairs  $(z, 1/z^*)$  by nearest partner  $\triangleright O(M^2)$ 
6: sort pairs by  $||z_{in}| - 1|$ ; keep the  $d$  smallest for  $k = 1, \dots, d$ 
   do
7:
   end
    $\hat{\theta}_k \leftarrow \arcsin(\arg z_k / \pi)$ ;  $\nu_k \leftarrow |\log |z_k||$   $\triangleright O(1)$ 
8:
9: flag  $\leftarrow [\max_k \nu_k > \tau]$ 
10: return  $\hat{\theta}$ ,  $\nu$ , flag = 0

```

angles $\theta_1, \dots, \theta_d$. With $\mu = \pi \sin \theta$ the manifold is $\mathbf{a}(\theta) = [1, e^{j\mu}, \dots, e^{j(M-1)\mu}]^T$ and $\mathbf{x}(t) = \mathbf{A}\mathbf{s}(t) + \mathbf{n}(t)$ with spatially white circular Gaussian noise. Let $\hat{\mathbf{R}} = N^{-1} \sum_t \mathbf{x}(t) \mathbf{x}^H(t)$, let $\hat{\mathbf{E}}_n$ span the estimated noise subspace and $\mathbf{C} = \hat{\mathbf{E}}_n \hat{\mathbf{E}}_n^H$. Root-MUSIC roots

$$p(z) = \sum_{l=-(M-1)}^{M-1} c_l z^{-l}, \quad c_l = \sum_{m-n=l} \mathbf{C}_{mn}, \quad (3)$$

retaining the d roots inside the unit circle closest to it and reporting $\hat{\theta}_k = \arcsin(\arg z_k / \pi)$.

Line 5 is the only step that differs from a textbook implementation, and it is $O(M^2)$ against the $O(M^3)$ of the rooting it follows. Line 10 is the detector, and costs one logarithm per source. Substituting $f_{\min,k} = \mathbf{a}^H(\hat{\theta}_k) \mathbf{C} \mathbf{a}(\hat{\theta}_k)$ for ν_k in line 10 gives the stronger detector of Section V-B at an added $O(dM^2)$.

We adopt the standard assumptions: the array is calibrated and uniform, the sources are narrowband and in the far field, the noise is spatially white circular Gaussian and independent of the signals, and $N > M$ so that $\hat{\mathbf{R}}$ is nonsingular almost surely. Sources are taken uncorrelated except in the coherent study of Table VIII, where spatial smoothing is applied first. The number of sources d is assumed known; Table VI includes an information-theoretic criterion only as a detection baseline, not as an enumeration step.

Two conventions matter for what follows. First, the root pairing in line 5 of Algorithm 2 is by nearest partner under $z \mapsto 1/z^*$, which is unambiguous whenever no two pairs are closer to each other than each is to its own partner; this can fail when two sources are unresolved, and is one route to the breakdown events studied in Section V. Second, we take the root strictly inside the unit circle as the representative of each pair, so $\nu_k = |\log |z_k|| = -\log |z_k| > 0$.

B. Pairing and double zeros

Proposition 1 (Conjugate-reciprocal pairing). *Since $\mathbf{C}$ is Hermitian, $c_{-l} = c_l^*$ and $p(z) = \overline{p(1/\bar{z})}$, so the roots of $z^{M-1} p(z)$ occur in pairs $(z, 1/z^*)$. Writing $z = e^{v+j\omega}$, the partner is*

$e^{-v+j\omega}$: the two share an argument and their geometric mean $e^{j\omega}$ lies exactly on the unit circle.

Proof: $\overline{p(1/\bar{z})} = \sum_l c_l^* z^l = \sum_l c_{-l} z^l = p(z)$, so $p(z_0) = 0$ implies $p(1/\bar{z}_0) = 0$. If $z = e^{v+j\omega}$ then $1/\bar{z} = e^{-v+j\omega}$, and their product is $e^{2j\omega}$. ■

Proposition 2 (Non-negativity and double zeros). *On the unit circle, $f(\omega)p(e^{j\omega}) = \mathbf{a}^H(\omega) \mathbf{C} \mathbf{a}(\omega) = \|\hat{\mathbf{E}}_n^H \mathbf{a}(\omega)\|^2 \geq 0$ for every realization. In the noiseless limit $f(\omega_k) = 0$ at each true angle, and since $f \geq 0$ these zeros are of even order, generically double. Consequently each signal pair $(z_k, 1/z_k^*)$ coalesces at $e^{j\omega_k}$ in that limit, a degenerate pair with spread zero, and splits radially under noise.*

Proof: Expanding, $\mathbf{a}^H \mathbf{C} \mathbf{a} = \sum_{m,n} e^{-jm\omega} \mathbf{C}_{mn} e^{jn\omega} = \sum_l c_l e^{-jl\omega} = f(\omega)$, which equals $\|\hat{\mathbf{E}}_n^H \mathbf{a}\|^2 \geq 0$. A non-negative smooth function vanishing at ω_k has $f'(\omega_k) = 0$, so the zero has order at least two. By Proposition 1 the two members of the pair have equal argument, so their coalescence point lies on the circle. ■

This is the structural heart of the paper. Root-MUSIC is not a generic rooting problem with well-separated roots; it is a problem in which d pairs are *exactly degenerate* in the noiseless limit, and the entire finite-sample behaviour of the signal roots is the splitting of those degenerate pairs. In the classification of [11], every signal pair sits on the boundary case $\Lambda = 0$.

C. An exact identity for the pair spread

Theorem 2 (Spread identity). *Let $\hat{\omega}_k$ be a local minimizer of f with value $f_{\min} = f(\hat{\omega}_k) > 0$ and curvature $f''(\hat{\omega}_k) > 0$, and let $z_k = e^{v_k+j\hat{\omega}_k}$ be the associated root inside the unit circle. Then*

$$\nu_k^2 = \frac{2f_{\min}}{f''(\hat{\omega}_k)} + O(\nu_k^4). \quad (4)$$

Proof: Let $g(v, \omega) = p(e^{v+j\omega}) = \sum_l c_l e^{-lv} e^{-jl\omega}$, so $g(0, \omega) = f(\omega)$. Differentiating, $\partial_v g|_{v=0} = -\sum_l l c_l e^{-jl\omega}$ while $f'(\omega) = -j \sum_l l c_l e^{-jl\omega}$, so $\partial_v g|_{v=0} = -j f'(\omega)$, which vanishes at a minimizer. Differentiating again, $\partial_v^2 g|_{v=0} = \sum_l l^2 c_l e^{-jl\omega}$ while $f''(\omega) = -\sum_l l^2 c_l e^{-jl\omega}$, so $\partial_v^2 g|_{v=0} = -f''(\omega)$. Taylor expanding in v at $\omega = \hat{\omega}_k$,

$$g(v, \hat{\omega}_k) = f_{\min} - \frac{1}{2} f''(\hat{\omega}_k) v^2 + O(v^4),$$

and setting $g = 0$, since z_k is a root, gives (4). ■

Identity (4) says the pair spread is the square root of the residual null depth normalized by curvature. It explains both why the spread is informative, since $f_{\min}$ measures how badly the estimated noise subspace fails to annihilate the steering vector, and, as Section V-B shows, why it is not the best available statistic.

Corollary 1 (Scaling). *Since the subspace estimation error satisfies $f_{\min} = O(1/(N \cdot \text{SNR}))$ under standard conditions [6], (4) gives $\nu_k = O((N \cdot \text{SNR})^{-1/2})$.*

Table III tests (4) directly, comparing the measured spread against the right-hand side evaluated by numerical differentiation of f . The median relative deviation is 0.19 to 0.28,

Table III: Test of the spread identity (4), $M = 8$, $d = 2$, sources at $\pm 5^\circ$, 250 trials per row.

SNR (dB)	N	median ν meas.	median ν pred.	rel. dev.
10	200	0.01405	0.01510	0.207
5	200	0.02506	0.02702	0.192
0	200	0.04644	0.05171	0.233
0	100	0.06552	0.07073	0.238
-5	200	0.09396	0.10298	0.282
-5	100	0.13026	0.14338	0.249

Table IV: Median pair spread versus snapshots, 200 trials per cell, and the fitted exponent against the predicted $N^{-1/2}$.

SNR	$N=50$	100	200	400	800	fitted
10	0.0282	0.0198	0.0138	0.0098	0.0070	$N^{-0.505}$
5	0.0503	0.0349	0.0250	0.0179	0.0127	$N^{-0.492}$
0	0.0933	0.0655	0.0455	0.0330	0.0237	$N^{-0.494}$
-5	0.1804	0.1303	0.0936	0.0669	0.0465	$N^{-0.487}$

consistent with the neglected $O(\nu^4)$ term. Table IV tests Corollary 1: the fitted exponents in N are -0.505 , -0.492 , -0.494 , -0.487 against the predicted $-1/2$, and the fitted exponent in linear SNR is -0.551 against $-1/2$.

IV. A NEGATIVE RESULT: SYMMETRIZATION GAINS NOTHING

Proposition 3 (Exact redundancy). *Let $(z_k, 1/z_k^*)$ be a root pair of (3). Then $\arg z_k = \arg(1/z_k^*)$ exactly, so the pair-symmetrized angle estimate $\frac{1}{2}[\arg z_k + \arg(1/z_k^*)]$ equals the textbook estimate $\arg z_k$ identically, for every realization and every sample size.*

Proof: Immediate from $z = e^{\nu+j\omega} \Rightarrow 1/\bar{z} = e^{-\nu+j\omega}$: the perturbation that moves a signal root off the unit circle is purely radial in the logarithmic coordinate and does not perturb the argument. ■

This deserves emphasis because it contradicts a plausible expectation. Symmetry exploitation often buys a variance reduction, as with forward-backward averaging [7]. Here it does not, because the symmetry is a property of the polynomial itself for *any* Hermitian $\mathbf{R}$, not a population property that the sample estimate inherits only approximately. There is no residual asymmetry to average away. Table V confirms the two estimators agree to all reported digits at every SNR, and that both attain the stochastic Cramér-Rao bound above threshold. The last column shows the correlation between $\log \nu_k$ and $\log |\hat{\theta}_k - \theta_k|$ is negligible above threshold: the spread is *not* a useful fine-grained reliability measure, and rises only in the breakdown region.

V. BREAKDOWN DETECTION

A. Setup

Subspace DOA estimators exhibit a threshold effect: below a critical SNR the estimate departs abruptly from the Cramér-Rao bound because estimated signal and noise subspaces swap a direction and a noise root is selected in place of a signal root [8], [9]. Detecting such events per snapshot block is

Table V: Root-MUSIC RMSE in degrees, $M = 8$, $d = 2$ at $\pm 5^\circ$, $N = 100$, 600 trials. “Pair” is the symmetrized estimator.

SNR (dB)	RMSE std	RMSE pair	CRB	$\rho(\log \nu, \log \text{err})$
-10	13.2856	13.2856	1.7313	0.338
-5	0.7666	0.7666	0.7522	0.048
0	0.3778	0.3778	0.3749	0.027
5	0.2018	0.2018	0.2015	0.039
10	0.1142	0.1142	0.1116	-0.020
15	0.0649	0.0649	0.0624	-0.030
20	0.0351	0.0351	0.0351	0.021

Table VI: Breakdown detection, area under the ROC curve. $M = 8$, $d = 2$, $N = 50$, 1400 trials per SNR.

SNR (dB)	$P(\text{brk})$	AUC ν	AUC $f_{\min}$	AUC gap	AUC AIC
-12	0.661	0.629	0.876	0.560	0.476
-10	0.491	0.722	0.956	0.611	0.469
-8	0.196	0.881	0.993	0.681	0.517
-6	0.029	0.974	1.000	0.842	0.575

practically valuable, since a flagged block can be discarded, down-weighted in a tracker, or re-estimated with more data.

Proposition 2 suggests a statistic. A signal root sits on the circle in the noiseless limit, so ν_k vanishes there; a noise root has no reason to be near the circle. We therefore test $\max_k |\nu_k|$, and by Theorem 2 also the closely related $f_{\min}$.

We take $M = 8$, $d = 2$ equal-power uncorrelated sources at $\pm 5^\circ$, $N = 50$, and declare breakdown when $\max_k |\hat{\theta}_k - \theta_k|$ exceeds one beamwidth, taken as $2/M$ radians or 14.32° . Baselines are the eigenvalue gap λ_d/λ_{d+1} and an AIC statistic $-2N(M-d) \log(\bar{g}/\bar{a}) + 2d(2M-d)$ with $\bar{g}, \bar{a}$ the geometric and arithmetic means of the noise eigenvalues [10].

B. Ablation

Table VI carries the central experimental finding, and it is not the one we anticipated. Both pair-derived statistics dominate the classical baselines by wide margins: at -6 dB the spread reaches AUC 0.974 against 0.842 for the eigenvalue gap and 0.575 for AIC, which is barely above chance because AIC is designed for source enumeration rather than breakdown. But the raw null depth $f_{\min}$ beats the spread at every SNR tested, reaching 1.000 at -6 dB and 0.956 at -10 dB where the spread manages only 0.722.

Theorem 2 explains this. Since $\nu_k^2 \approx 2f_{\min}/f''$, the spread is the null depth divided by a curvature that itself fluctuates from trial to trial. That division injects variance without adding information about the event being detected, and degrades the statistic. The identity we derived to justify the spread turns out to predict that the spread should be inferior, and the measurement agrees.

The practical consequence is a genuine trade. The spread is free: the moduli $|z_k|$ are produced by the same rooting step that produces the angles, so ν_k costs one logarithm per source, measured at $1.2 \mu\text{s}$ independent of M (Table X). Computing $f_{\min}$ requires one additional evaluation of $\mathbf{a}^H \mathbf{C} \mathbf{a}$ at each estimated angle, an $O(M^2)$ operation. A practitioner already rooting and wanting a detector at no cost should use ν ; one able to afford $O(dM^2)$ extra should use $f_{\min}$.

Table VII: Detection probability at fixed false-alarm rate. $M = 8$, $d = 2$, $N = 50$, 1600 trials per SNR; thresholds calibrated on non-breakdown trials.

SNR	$P(\text{brk})$	P_d using ν			P_d using $f_{\min}$		
		0.01	0.05	0.10	0.01	0.05	0.10
-12	0.674	0.023	0.141	0.235	0.051	0.331	0.596
-10	0.486	0.093	0.204	0.316	0.265	0.712	0.919
-8	0.199	0.148	0.425	0.607	0.877	1.000	1.000
-6	0.022	0.500	0.861	0.972	1.000	1.000	1.000

Table VIII: Two fully coherent sources at $\pm 6^\circ$, $M = 10$, forward spatial smoothing with $L = 3$, $N = 80$, 900 trials per SNR.

SNR (dB)	$P(\text{brk})$	RMSE (deg)	AUC ν	AUC gap
-6	0.051	9.8601	0.756	0.646
-3	0.000	1.2534	—	—
0	0.000	0.7247	—	—
3	0.000	0.4501	—	—

C. Operating points at fixed false-alarm rate

Area under the ROC curve summarizes a detector over all thresholds; a practitioner needs performance at a threshold they can actually set. Table VII therefore reports detection probability at three false-alarm rates, with the threshold calibrated on the non-breakdown trials at each SNR.

The ordering of Section V-B persists and sharpens. At -8 dB and a five percent false-alarm rate, $f_{\min}$ detects every breakdown in 1600 trials while ν detects 42.5%. Deep below threshold both degrade, and at -12 dB neither is usable at a one percent false-alarm rate. The usable operating region for a free detector is therefore the upper part of the transition, roughly -8 to -6 dB in this configuration, which is where breakdown is intermittent and detection is most valuable; below that, the correct response is to collect more data rather than to flag blocks.

D. Coherent sources

Multipath produces coherent arrivals, for which the signal covariance is rank deficient and MUSIC fails without preprocessing. Spatial smoothing [7] restores rank by averaging subarray covariances, and because the smoothed covariance is again Hermitian, Proposition 1 continues to hold: the pair structure survives preprocessing unchanged.

Table VIII reports a 10-sensor array with two fully coherent sources at $\pm 6^\circ$, forward smoothing with $L = 3$ subarrays of size 8, and 80 snapshots. Estimation is accurate once above threshold, and the spread retains its advantage over the eigenvalue gap in the breakdown region, 0.756 against 0.646 at -6 dB. The margin is smaller than in the uncorrelated case, which is expected: smoothing reduces the effective aperture and flattens the null spectrum, so by Theorem 2 the curvature f'' shrinks and the spread becomes noisier.

E. Sensitivity to array size and separation

Table IX varies aperture, separation and SNR. The spread retains a clear advantage over the eigenvalue gap in the

Table IX: Detector sensitivity, $N = 60$, 1200 trials per row, two sources.

M	sep. (deg)	SNR (dB)	$P(\text{brk})$	AUC ν	AUC gap
6	10	-6	0.296	0.813	0.682
8	10	-6	0.017	0.989	0.848
8	6	-4	0.263	0.882	0.715
8	14	-8	0.004	0.966	0.986
10	8	-6	0.001	0.989	0.972

Table X: Measured cost per block in microseconds, 40 repetitions.

M	deg	eigendecomp.	rooting	spread
6	10	11.8	55.9	1.23
8	14	13.0	100.5	1.27
12	22	25.4	225.0	1.11
16	30	38.8	433.2	1.17
24	46	90.2	995.2	1.21
32	62	109.8	2011.3	1.22

small-aperture and closely-spaced regimes, where breakdown is common: at $M = 6$ with 10° separation the AUC is 0.813 against 0.682, and at $M = 8$ with 6° separation it is 0.882 against 0.715. The advantage narrows and can reverse when breakdown becomes rare: at $M = 8$ with 14° separation the gap statistic slightly leads, 0.986 against 0.966. We report this rather than restricting the sweep.

At $M = 12$ with 10° separation no breakdown occurred in 1200 trials, so no AUC is defined; this is the regime in which the detector is unnecessary.

F. Complexity

Proposition 4 (Cost). *Per snapshot block, root-MUSIC costs $O(M^3)$ for the eigendecomposition, $O(M^2)$ to form the coefficients (3), and $O(M^3)$ to root the degree-2($M-1$) polynomial by companion-matrix eigenvalues. The spread adds $O(d)$; the null depth adds $O(dM^2)$.*

Table X confirms that rooting dominates, growing by a factor of 36 from $M = 6$ to $M = 32$ against the factor ≈ 150 a pure cubic in the polynomial degree would suggest, while the spread stays at about $1.2 \mu\text{s}$ throughout. Relative to the 2011 μs rooting cost at $M = 32$, the detector is free to within measurement noise.

G. Implementation notes

Three practical points arose in our implementation and are worth recording.

The polynomial $z^{M-1}p(z)$ has degree $2(M-1)$ and its coefficients decay rapidly away from the centre for well-separated sources, so forming it in the natural order and rooting by companion-matrix eigenvalues is adequate up to the array sizes in Table X. For very large apertures a balanced companion form would be advisable, since the coefficient dynamic range grows with M .

The pair grouping should be done by matching each root to its nearest partner under $z \mapsto 1/z^*$ and removing both from the pool, rather than by simply taking all roots inside the unit

circle. The two agree when the pairing is unambiguous, but the former degrades more gracefully when two pairs nearly collide, because it never assigns two inside-roots to the same partner.

When a signal root is very close to the unit circle, $|z_k| \rightarrow 1$ and $\nu_k \rightarrow 0$, the logarithm is well conditioned, but the *pairing* is not: the two members of the pair coalesce, which is precisely the degenerate configuration of Proposition 2. In that regime the two roots returned by the eigenvalue solver may differ only at the level of rounding error, and ν_k inherits that error. This is harmless for detection, since the regime corresponds to a well-estimated source, but it means ν_k should not be interpreted quantitatively below roughly $\sqrt{\epsilon_{\text{mach}}}$.

VI. RELATED WORK AND POSITIONING

Root-MUSIC and its analysis. Root-MUSIC originates with Barabell [4]; the definitive finite-sample analysis is Rao and Hari [5], who established asymptotic efficiency for uncorrelated sources on a uniform linear array, and Stoica and Nehorai [6] for the Cramér-Rao comparison. Our Table V reproduces that agreement and we make no claim of improving on it. The contribution here is orthogonal: it concerns a statistic that the standard algorithm computes and discards.

Performance breakdown. The threshold effect in subspace methods was analysed as a subspace swap by Thomas, Scharf and Tufts [8], and Johnson, Abramovich and Mestre [9] compared breakdown across MUSIC, G-MUSIC and maximum likelihood. That literature is largely predictive, characterizing when breakdown occurs as a function of SNR and aperture. The detectors studied here are instead *per-realization*: they aim to flag the particular block that failed, which is what a tracker needs. To our knowledge the null-spectrum depth at the estimated angle has not been evaluated as a breakdown statistic, and the ordering we establish between it and the radial spread is new, as is the identity of Theorem 2 that predicts that ordering.

Source enumeration. Information-theoretic criteria [10] solve a different problem, choosing d rather than assessing a fixed- d estimate. Table VI includes AIC only to confirm that it is near chance for this task, which is expected rather than a criticism of the criterion.

Structure exploitation. Forward-backward averaging [7] is the canonical example of a symmetry that does reduce variance. Proposition 3 shows that the conjugate-reciprocal symmetry is not such an example, and the distinction is instructive: forward-backward averaging exploits a symmetry of the *population* covariance that the sample estimate violates, leaving a residual to average away, whereas the conjugate-reciprocal symmetry holds identically for every sample covariance and leaves no residual.

Cascade realization. Second-order-section implementation and the Schur triangle are standard [1], [2]. What Section II-B adds is a controlled measurement of how much the parameterization matters during *optimization*, as opposed to during implementation, and the answer, 40/40 against 1/40 stable designs, is larger than we expected.

VII. DISCUSSION: WHAT THE IDENTITY BUYS

It is worth separating what Theorem 2 contributes from what the experiments contribute, because the two play different roles.

The experiments alone would establish an ordering, that $f_{\min}$ detects breakdown better than ν , which detects it better than the eigenvalue gap. An ordering measured at four SNRs in one array configuration is weak evidence for a general recommendation: it could easily be an artifact of the geometry, the snapshot count, or the breakdown labelling. The identity is what upgrades it. Because $\nu^2 \approx 2f_{\min}/f''$ holds for *any* Hermitian covariance estimate and any array for which p is conjugate-reciprocal, the relation between the two statistics is structural. The spread is the null depth divided by a quantity that carries no information about whether this particular block broke down, but which fluctuates across blocks. A statistic formed by dividing an informative quantity by an independent noisy one is necessarily the weaker of the two, and the ordering follows without appeal to the experiment.

The identity also predicts *where* the gap should be largest. The curvature f'' measures the sharpness of the null at the estimated angle, which degrades as aperture shrinks, as sources approach, and as smoothing is applied. In each of those regimes ν should lose ground to $f_{\min}$, and Tables IX and VIII are consistent with that: the spread's margin over the gap statistic is largest at $M = 6$ and at 6° separation, where breakdown is frequent, and the coherent case with its reduced effective aperture shows the narrowest margin of all the regimes in which breakdown occurs.

Finally, the identity explains a fact that would otherwise look like a puzzle. Table V shows almost no correlation between the spread and the angular error above threshold, yet Table VI shows the spread detecting breakdown very well. These are consistent because $f_{\min}$ measures a *global* property, how far the estimated noise subspace is from annihilating the steering vector, while the angular error is driven by the component of the subspace perturbation along a particular direction. Above threshold the two are nearly independent; at breakdown the global quantity degrades sharply and the angular error follows.

VIII. LIMITATIONS

Several limitations should be stated. The spread identity (4) is first order; Table III shows a systematic relative deviation of 0.19 to 0.28, growing as SNR falls, which is the neglected quartic term. It is adequate for predicting scaling and for explaining the detector ordering, but it is not a high-accuracy predictor of ν in the threshold region, which is exactly where the detector operates.

The detection results are empirical ROC curves, not a test with a guaranteed false-alarm rate. Setting a threshold in deployment requires the null distribution of $\max_k f_{\min,k}$, which we have not derived; the thresholds in Table VII were calibrated on non-breakdown trials, which presumes a labelled calibration set.

Our breakdown label is a fixed multiple of the beamwidth. Alternative definitions, such as departure from the Cramér-Rao

bound by a factor, would shift the labels and hence the AUC values, although we would not expect the ordering between statistics to change, since it follows from Theorem 2 rather than from the labelling convention.

Finally, all array results assume a calibrated uniform linear array with uncorrelated or fully coherent sources and spatially white noise. Array imperfections perturb the conjugate-reciprocal symmetry, and how gracefully the pair grouping in line 5 of Algorithm 2 degrades under miscalibration is not addressed here.

IX. CONCLUSION

Respecting the pair structure of the polynomials arising in cascade realization and in subspace DOA estimation has sharply different consequences in the two settings. For cascade design, pair coordinates make the Schur region a triangle per section and a polytope for the cascade, turning constrained magnitude fitting into a smooth problem over a convex set, with near-geometric error decay in the number of sections and the spread sign reporting the character of each section for free.

For root-MUSIC the picture is more interesting and less flattering to our initial hypothesis. The signal pairs are exactly degenerate in the noiseless limit, so the finite-sample problem is the splitting of double roots, and the splitting obeys the identity $v^2 = 2f_{\min}/f''$, verified here along with its $N^{-1/2}$ consequence. Symmetrizing over a pair gains nothing whatsoever, which we have proved rather than merely observed. The spread detects breakdown far better than the eigenvalue gap or AIC, but worse than the raw null depth, and our own identity predicts exactly that ordering. The defensible recommendation is therefore conditional on budget rather than uniform, and we prefer to state it that way.

Natural extensions are to non-uniform and sparse arrays, where the rooting polynomial keeps its conjugate-reciprocal symmetry but the selection rule changes; to coherent sources after spatial smoothing, where the symmetry survives; and to a sequential detector accumulating $f_{\min}$ across blocks inside a tracker. A theoretical question left open is the distribution of $f_{\min}$ under the breakdown hypothesis, which would turn the ROC curves reported here into a threshold with a guaranteed false-alarm rate.

REFERENCES

- [1] A. V. Oppenheim and R. W. Schaffer, *Discrete-Time Signal Processing*, 3rd ed. Upper Saddle River, NJ: Prentice Hall, 2009.
- [2] S. K. Mitra, *Digital Signal Processing: A Computer-Based Approach*, 4th ed. New York: McGraw-Hill, 2011.
- [3] R. O. Schmidt, "Multiple emitter location and signal parameter estimation," *IEEE Trans. Antennas Propag.*, vol. 34, no. 3, pp. 276–280, 1986.
- [4] A. J. Barabell, "Improving the resolution performance of eigenstructure-based direction-finding algorithms," in *Proc. IEEE ICASSP*, 1983, pp. 336–339.
- [5] B. D. Rao and K. V. S. Hari, "Performance analysis of root-MUSIC," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 37, no. 12, pp. 1939–1949, 1989.
- [6] P. Stoica and A. Nehorai, "MUSIC, maximum likelihood, and Cramér-Rao bound," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 37, no. 5, pp. 720–741, 1989.
- [7] S. U. Pillai and B. H. Kwon, "Forward/backward spatial smoothing techniques for coherent signal identification," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 37, no. 1, pp. 8–15, 1989.

- [8] J. K. Thomas, L. L. Scharf, and D. W. Tufts, "The probability of a subspace swap in the SVD," *IEEE Trans. Signal Process.*, vol. 43, no. 3, pp. 730–736, 1995.
- [9] B. A. Johnson, Y. I. Abramovich, and X. Mestre, "MUSIC, G-MUSIC, and maximum-likelihood performance breakdown," *IEEE Trans. Signal Process.*, vol. 56, no. 8, pp. 3944–3958, 2008.
- [10] M. Wax and T. Kailath, "Detection of signals by information theoretic criteria," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 33, no. 2, pp. 387–392, 1985.
- [11] S. Pratiher, "On the average-offset pair-partitioning of polynomial roots," submitted, 2026 (<https://zenodo.org/records/19411173>).