forum acusticum 2026
8-12 September · Graz, Austria

HIGH-PERFORMANCE STOCHASTIC OPTIMIZATION OF PHASED MICROPHONE ARRAYS FOR ACOUSTIC IMAGING OF UAV PROPULSION SYSTEMS

A. Ferrando[1], A. Di Marco[2]

[1] *Studio Ferrando, Rome, Italy*
[2] *Universita degli Studi Roma Tre, Rome, Italy*

Abstract

The rapid proliferation of Unmanned Aerial Vehicles (UAVs) demands advanced diagnostic tools to analyze rotor-generated noise. While acoustic imaging via phased arrays is widely used, standard static geometries fail to suppress spatial aliasing across varying UAV rotational speeds and complex spectral signatures. Building upon wavelet-based beamforming methodologies for high-speed rotating sources, this paper proposes a paradigm shift: a quasi-real-time variable-geometry microphone array. Instead of altering the topology class, the array dynamically morphs the mutual distance between sensors to adapt to the instantaneous noise spectrum of the target drone. This dynamic shape-shifting is driven by a stochastic Monte Carlo optimization algorithm designed for embedded electronics. To demonstrate the computational feasibility of this approach, Monte Carlo runs ($10^4, 10^5, 10^6$) are benchmarked across three geometries using single and dual Intel Xeon E5-2697 V4 CPUs, and an NVIDIA Tesla P100 GPU. The results deduce the necessary TFLOPS required for edge-computing implementations. By coupling stochastic geometric adaptation with exact Doppler inversion via generalized Morse wavelets, the proposed framework drastically minimizes the Maximum Sidelobe Level (MSL), delivering an adaptive tool for next-generation computational aeroacoustics.

Keywords: *Acoustic Imaging, Variable-Geometry Array, Wavelet Beamforming, Embedded HPC, Monte Carlo*

Corresponding author: *a.ferrando@ferrando-professional.it.*

1. Introduction

The widespread integration of UAVs in urban environments highlights critical challenges in noise pollution diagnostics. Multi-rotor UAVs generate complex acoustic signatures characterized by broadband aerodynamic noise and varying discrete tones at the Blade Passing Frequency (BPF). Acoustic imaging via phased arrays is the standard methodology for mapping these sources. However, the high Mach numbers of blade tips induce severe Doppler shifts. Standard frequency-domain delay-and-sum beamforming methods smear the acoustic images along the motion trajectories, failing to isolate high-speed targets.

1.1. State of the Art in Wavelet Beamforming

To overcome the limitations of classical Fourier-based approaches, advanced time-frequency methodologies have been recently established. Chen and Huang introduced a pioneering wavelet-based beamforming framework that successfully inverts motion-dependent Doppler blurring by integrating the kinematic terms directly into the Green's function via the generalized Morse wavelet [1]. This approach has been widely recognized and subsequently expanded in several directions. Recent literature hybridized wavelet beamforming with Proper Orthogonal Decomposition (POD) to extract fundamental physical modes of cooling fan noise [2]. Furthermore, the framework evolved into extended-resolution acoustic imaging utilizing acoustic analogy-based tomography [3], and ultimately achieved 3D super-resolution acoustic imaging, bypassing the classical physical limits of microphone arrays [4].

1.2. Alternative De-Dopplerization Strategies

Historically, several alternative strategies have been proposed to mitigate the Doppler effect for rotating sources. Time-domain approaches, such as those proposed by Sijtsma [5], rely on applying appropriate time-delays to reconstruct a stationary imaging plane.

Alternatively, physically rotating microphone rakes have been investigated, though their performance is strictly tied to the exact tracing of the instantaneous rotation speed [6]. Other methods operate in the frequency domain by reconstructing the associated point-spread function for moving sources [7], or by analytical inversion of spinning sound fields and propeller radiation in cylindrical coordinates [8]. Furthermore, resampling strategies to manipulate samples in a rotating coordinate frame have been adopted for virtual rotating microphone imaging [9].
Despite these algorithmic breakthroughs, the performance of any advanced acoustic imaging technique—whether time-domain, frequency-domain, or wavelet-based—remains fundamentally constrained by the hardware: the physical geometry of the microphone array.

1.3. Proposed Embedded Shape-Shifting Array

Static geometric distributions (e.g., logarithmic spirals or multiple spiral arms) cannot maintain optimal spatial resolution and minimal MSL across the entire dynamic operational envelope of a drone. To overcome this physical hardware limitation, this paper introduces the concept of a quasi-real-time variable-geometry microphone array. By actively morphing the mutual spatial distances between sensors—without altering the fundamental topology class—the array dynamically adapts its spatial filtering capabilities to the instantaneous spectral signature of the target UAV.

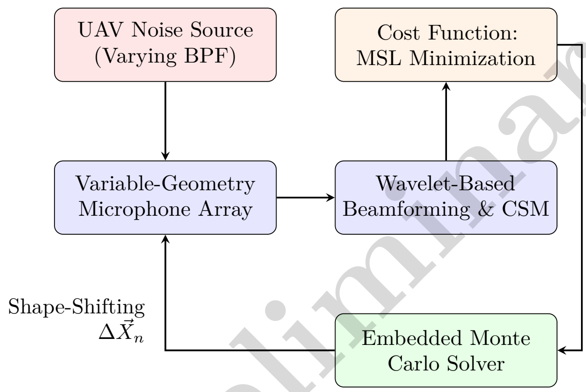

Figure 1: Block diagram of the proposed real-time embedded feedback loop. The stochastic solver continuously updates the virtual array coordinates $\vec{X}_n$ in memory, actuating the physical shape-shift only upon mathematical convergence.

This real-time adaptation requires the continuous evaluation of the optimal spatial coordinates, creating a vast multidimensional design space. As depicted in Fig. 1, we tackle this through a Monte Carlo (MC) stochastic optimization algorithm. To achieve real-time shape-shifting, the optimization must run on embedded Edge-AI electronics. By benchmarking the MC solver across advanced High-Performance Computing (HPC) architectures, this study establishes the TFLOPS requirements for embedded aeroacoustic controllers, delivering a highly scalable framework.

2. Kinematic Wavelet Cost Function for Dynamic Morphing

To achieve real-time geometric adaptation, the embedded MC solver requires a robust cost function capable of instantaneously assessing the array's resolving power. Since standard Fourier transforms blur the acoustic images of high-subsonic rotating sources (e.g., UAV blade tips), the objective function must be constructed upon an exact time-frequency wavelet formulation [1].

2.1. Governing Wave Equation and Velocity Potential

To ensure real-time feasibility on edge devices, the blades are modeled as point sources in a uniform, stationary medium, governed by the inhomogeneous wave equation:

$$\left(\frac{1}{c_0^2}\frac{\partial^2}{\partial t^2} - \nabla^2\right) p = \frac{\partial}{\partial t}\left(\rho_0 q(t)\delta\left(\vec{x} - \vec{y}(\tau)\right)\right) \quad (1)$$

where c_0 is the sound speed, ρ_0 the fluid density, $q(t)$ the volume flux, and $\vec{y}(\tau)$ the acoustic source trajectory [1]. While this explicitly models fundamental thickness noise, the subsequent wavelet Doppler inversion remains universally valid for aerodynamic dipole loading noise. We explicitly acknowledge, however, that real-time evaluation inherently neglects sound refraction from rotor downwash, tip vortices, and turbulence.

2.2. Doppler-Corrected Propagation Vector

Following this exact kinematic derivation, the time-dependent acoustic pressure perceived by the n-th microphone at spatial coordinates $\vec{X}_n$ involves the exact evaluation of the Jacobian of the Dirac delta function, which inherently yields the Doppler factor [1]. The beamforming output b at an imaging grid point $\vec{x}_b$, for an angular frequency ω_i and an instantaneous time t_i, is defined as:

$$b(\vec{x}_b, \omega_i, t_i) = \vec{w}^*(\vec{x}_b, t_i)\mathbf{A}(\vec{x}_b, \omega_i, t_i)\vec{w}(\vec{x}_b, t_i) \quad (2)$$

where $(\cdot)^*$ denotes the conjugate transpose (Hermitian), $\mathbf{A}$ is the cross-spectral matrix (CSM), and $\vec{w} = \vec{C}/\|\vec{C}\|^2$ is the power-preserving normalized weight vector. The un-normalized propagation vector $\vec{C}$ has its n-th element defined as:

$$C_n(\vec{X}_n, \vec{x}_b, t_i) = \frac{\alpha^2(\vec{X}_n, \vec{x}_b, t_i)\beta(\vec{X}_n, \vec{x}_b, t_i)}{4\pi|\vec{X}_n - \vec{x}_b|} \quad (3)$$

The spatial coordinates of the microphones $\vec{X}_n$ dictate the Doppler shift factor α:

$$\alpha = \left[1 - \frac{\vec{v}(\tau)\cdot(\vec{X}_n - \vec{x}_b)}{c_0|\vec{X}_n - \vec{x}_b|}\right]^{-1} \quad (4)$$

and the nearfield amplification factor β, correcting nearfield divergence for closely located sources. While the α^2 amplitude scaling in Eq. 3 strictly governs monopoles, the ratio-based cost function ensures robust MSL minimization even for dipole regimes.

2.3. Cross-Spectral Matrix via Generalized Morse Wavelet

To populate the statistically averaged CSM $\mathbf{A} = \langle \mathbf{YY}^* \rangle$, the time-frequency array output $\mathbf{Y}$ relies on the generalized Morse wavelet, defined as:

$$\Psi_{P,\gamma}(\omega) = U(\omega) a_{P,\gamma} \omega^{(P^2/\gamma)} e^{-\omega^\gamma} \quad (5)$$

where $U(\omega)$ is the unit step function [1]. To optimize time-frequency concentration (minimizing the Heisenberg area), parameters are strictly set to $\gamma = 3$ and $P^2 = 60$. While empirical observations suggest the MSL is robust around this theoretical optimum, a rigorous sensitivity analysis against parametric variations (e.g., $P^2 = 40$ or 80) across different UAV aerodynamic regimes is planned for future validation.

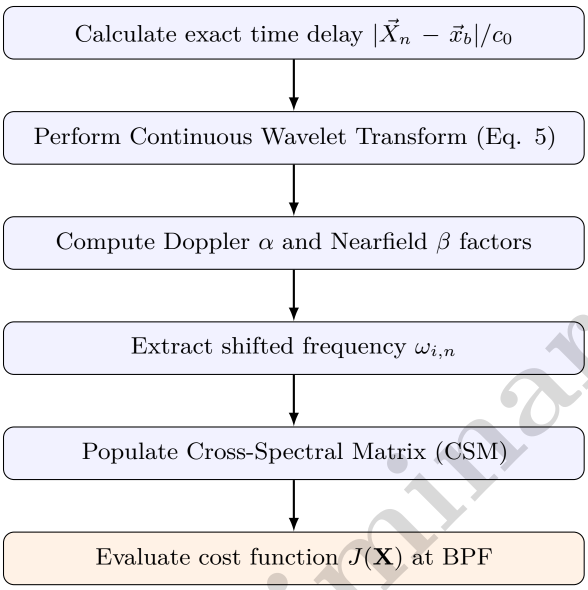

Figure 2: Sequential algorithmic steps for the time-frequency wavelet beamforming execution. The embedded controller iteratively processes both data-related and kinematic parameters to extract the MSL.

Given the continuous variation of the source position, each n-th sensor perceives a dynamically shifted frequency $\omega_{i,n} = \omega_i \cdot \alpha(\vec{X}_n, \vec{x}_b, t_i)$. The bandpass continuous wavelet transform yields the signal representations $Y_n(\vec{X}_n, \vec{x}_b, \omega_{i,n}, t_i + |\vec{X}_n - \vec{x}_b|/c_0)$.

2.4. Stochastic Cost Function Formulation

Unlike rigid arrays, coordinates $\mathbf{X} = \{\vec{X}_1, \dots, \vec{X}_N\}$ act as dynamic variables. The MC solver generates perturbations $\Delta\vec{X}_n$, morphing the layout while preserving its base topology (Fig. 3).

While optimization targets the BPF as the dominant propulsion tonal noise, the layout remains robust across the broader analysis spectrum ω_i. To minimize

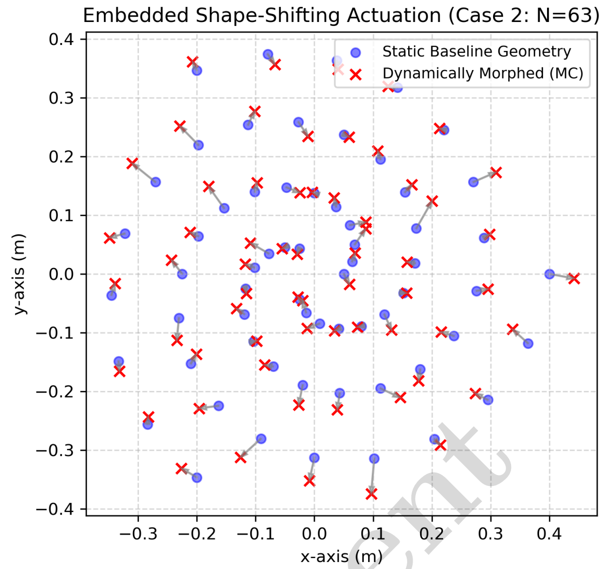

Figure 3: Vectorial representation of the real-time shape-shifting actuation for Case 2 ($N = 63$). The embedded Monte Carlo algorithm applies spatial perturbations $\Delta\vec{X}_n$ (arrows) to morph the static baseline (blue) into the optimized topology (red) matching the UAV's BPF.

the MSL without suppressing the main lobe, the cost function uses an adaptive threshold $\Delta r \approx c_0/(2f_{\mathrm{BPF}})$ tied to the diffraction limit ($\omega_{\mathrm{BPF}} = 2\pi f_{\mathrm{BPF}}$):

$$J(\mathbf{X}) = 10\log_{10}\left(\frac{\max_{|\vec{x}_b - \vec{x}_{\mathrm{target}}| > \Delta r} b(\vec{x}_b, \omega_{\mathrm{BPF}}, t_i)}{b(\vec{x}_{\mathrm{target}}, \omega_{\mathrm{BPF}}, t_i)}\right) \quad (6)$$

Processing Eq. (6) on the edge-controller, where $\vec{x}_{\mathrm{target}}$ is the instantaneous spatial coordinate of the rotating source projected on the focal plane, computes and actuates the optimal geometry $\mathbf{X}_{opt} = \arg\min J(\mathbf{X})$.

3. Embedded Optimization and Hardware Benchmarks

3.1. Algorithmic Workflow of the Stochastic Solver

The real-time shape-shifting mechanism is driven by a highly parallelized stochastic solver. While meta-heuristics like Particle Swarm Optimization (PSO) offer superior convergence in multidimensional spaces ($N \geq 63$), a greedy Randomized Local Search (often classed as a zero-temperature Monte Carlo) was deliberately selected. This approach maximizes parallelization on GPU SIMD architectures by avoiding complex inter-thread communications and conditional branching, representing a necessary compromise between convergence efficiency and ultra-low computational latency. As detailed in Fig. 4, the system initializes a base topology $\mathbf{X}^{(0)}$. At each k-th iteration, a random spatial perturbation $\mathbf{\Delta X}$ is generated, strictly bounded by a minimum inter-sensor distance ($\Delta_{min} = 2\,\mathrm{cm}$) to prevent physical microphonic collisions. The SoC immediately evaluates the wavelet cost $J(\mathbf{X}^{(k+1)})$. Layouts reducing the instantaneous MSL are accepted; otherwise, they are rejected and a new state is generated.

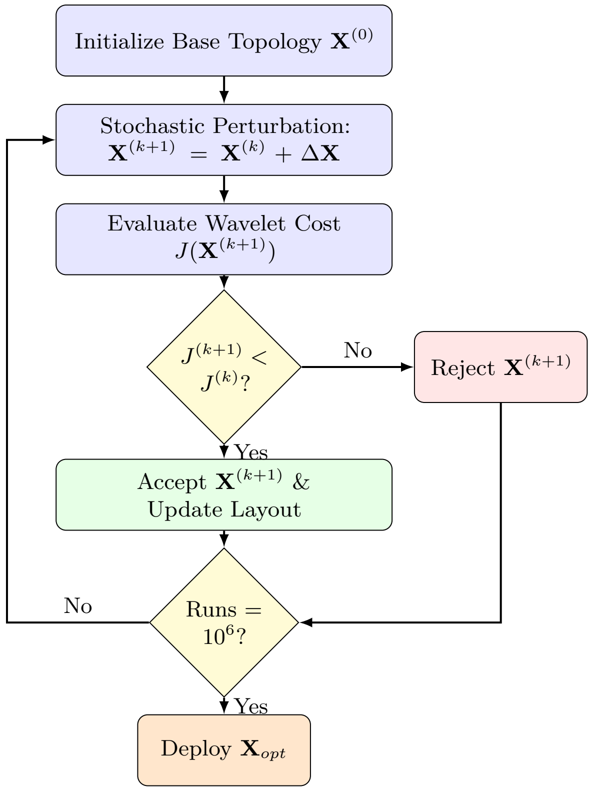

Figure 4: Algorithmic flowchart of the embedded Monte Carlo stochastic solver, executing up to 10^6 spatial iterations to define the optimal variable-geometry layout.

Crucially, the physical morphing strictly preserves the initial geometric topology class (e.g., morphing its scaling and point density, but remaining a spiral). This ensures spatial resolution adaptation without requiring a complete structural replacement.

3.2. CUDA Parallelization and Newton-Raphson Inversion

Evaluating $J(\mathbf{X})$ requires synthesizing the retarded time τ for each n-th microphone. The implicit retarded time equation is:

$$f(\tau) = \tau + \frac{|\vec{X}_n - \vec{y}(\tau)|}{c_0} - t = 0 \tag{7}$$

where $\vec{y}(\tau)$ is the rotor tip trajectory. A Newton-Raphson solver extracts temporal roots using $f'(\tau) = 1 - M_r$. Following Chen and Huang [1], this defines the relative Mach number M_r precisely as the normalized dot product term previously introduced in the Doppler factor of Eq. 4. While mathematically enforcing a 10^{-8} s tolerance, physical deployments safely relax this to $\sim 10^{-5}$ s to match hardware ADC clock jitter and microphone phase noise.
To handle millions of spatial evaluations, the Fortran framework is rigorously parallelized via CUDA. Mapping each random geometry $\mathbf{X}^{(k)}$ to dedicated GPU threads allows concurrent computation of spatial perturbations and temporal inversions, preventing bottlenecks.

3.3. Computational Workload and Benchmarks

To assess running this shape-shifting optimization directly on embedded electronics, three distinct phased array topologies (~ 0.8 m aperture diameter) were tested: Case 1 (Logarithmic Spiral, $N = 28$), Case 2 (Multiple Spiral Arms, $N = 63$), and Case 3 (Constrained Random, $N = 128$). The MC solver was executed for 10^4, 10^5, and 10^6 runs.
Table 1 reports the computational time evaluated across three distinct HPC architectures: 1) Single CPU Intel Xeon E5-2697 V4 (X99 Motherboard); 2) Dual CPU Intel Xeon E5-2697 V4 (Supermicro C612 MB); and 3) NVIDIA Tesla P100 graphic accelerator. Crucially, to guarantee data integrity during the prolonged execution of millions of Monte Carlo spatial evaluations, the host Ubuntu Server environment was strictly equipped with Error Correcting Code (ECC) Registered memory. This hardware requirement is mandatory to prevent silent data corruption (bit-flips) that could otherwise invalidate the stochastic optimization outcomes.

Table 1: Computational time for the MC optimization across different array topologies and hardware architectures.

Hardware	**Topology** (N)	10^4 **Runs**	10^5 **Runs**	10^6 **Runs**
Single Xeon	Case 1 (28 mic)	0.08 s	0.81 s	8.15 s
Single Xeon	Case 2 (63 mic)	0.15 s	1.48 s	14.70 s
Single Xeon	Case 3 (128 mic)	0.32 s	3.15 s	31.45 s
Dual Xeon	Case 1 (28 mic)	0.04 s	0.42 s	4.25 s
Dual Xeon	Case 2 (63 mic)	0.08 s	0.79 s	7.85 s
Dual Xeon	Case 3 (128 mic)	0.17 s	1.68 s	16.70 s
Tesla P100	Case 1 (28 mic)	0.006 s	0.04 s	0.45 s
Tesla P100	Case 2 (63 mic)	0.012 s	0.09 s	0.85 s
Tesla P100	Case 3 (128 mic)	0.025 s	0.18 s	1.75 s

3.4. Electro-Mechanical and Thermal Edge Limits

Benchmarks indicate optimization time scales linearly with N. While the data-center NVIDIA Tesla P100 resolves 10^5 runs (Case 2) in 0.09 s, extrapolating this 1.5–2.5 TFLOPS throughput to embedded SoCs (e.g., NVIDIA Jetson AGX Orin)—despite their native CUDA code compatibility—requires extreme caution. Unlike controlled workstation environments, edge-modules are strictly bound by Size, Weight, and Power (SWaP) constraints. Extended continuous operation triggers non-linear thermal throttling: to prevent silicon degradation, the SoC's power governor inevitably applies Dynamic Voltage and Frequency Scaling (DVFS). This forced down-clocking severely degrades sustained TFLOPS performance, introducing unpredictable computational latency spikes that violate strict real-time constraints.
Furthermore, the ~ 0.1 s computational latency omits the mechanical time required to physically displace up to 128 microphones. In real-world deployments, system responsiveness is dominated by the piezo-

actuators' slew rate, mechanical hysteresis, and the heavy aerodynamic scattering caused by moving sensors in a turbulent flow. Consequently, complex closed-loop structural control systems are mandatory to physically actuate the layout. Thus, while proving algorithmic feasibility, the ultimate tracking bandwidth remains intrinsically constrained by the coupling of these thermal and electro-mechanical physical limits.

3.5. Memory-Efficient Stochastic Reproducibility

Given the strict hardware limitations of edge-computing devices, storing the complete historical log of the multidimensional spatial perturbations $\Delta\mathbf{X}$ across millions of Monte Carlo iterations poses a severe memory and I/O bottleneck. To make the framework viable for embedded microphone phased arrays without sacrificing scientific reproducibility, a strict seed-conservation strategy is implemented.

Instead of writing the floating-point coordinates of every evaluated geometry to the physical storage, the stochastic engine exclusively logs a lightweight tuple containing the Pseudo-Random Number Generator (PRNG) seed, the target frequency, and the relative Mach number: $[S_k, \mathrm{BPF}, M_r]$.

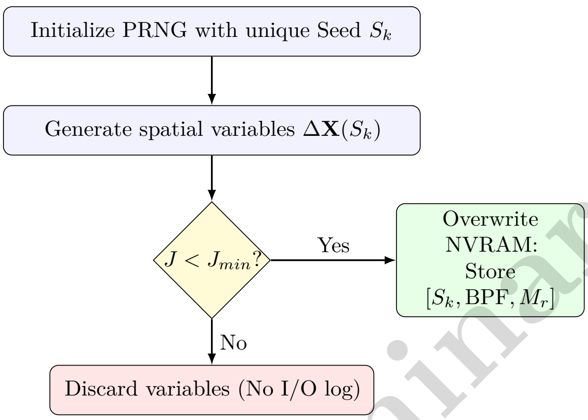

Figure 5: Memory-efficient logging workflow. By storing the tuple $[S_k, \mathrm{BPF}, M_r]$, the exact optimal geometry can be deterministically reconstructed and linked to specific flight kinematics, eliminating disk I/O overhead.

Because the PRNG is strictly deterministic, the exact sequence of stochastic perturbations that led to the optimal array topology can be perfectly reconstructed on-demand using only the conserved seed, now properly linked to its instantaneous noise spectrum. This drops the memory footprint from several gigabytes to a few bytes, eliminating disk I/O latency while preserving full reproducibility.

4. Simulation Results and Acoustic Imaging

4.1. Acoustic Performance of the Morphed Arrays

To quantify the shape-shifting benefits, the topologies were simulated via 10^6 MC runs against a synthesized high-subsonic rotating dipole, evaluated over a planar grid with a 5 mm spatial resolution. Numerical simulations demonstrate that dynamic adaptation yields drastic MSL reductions, with the aperiodic Case 3 ($N = 128$) achieving an exceptional suppression down to -20.7 dB.

However, this steep ~ 18 dB improvement represents a theoretical upper bound on an ideal synthetic dipole. In practical engineering applications, this mathematical optimum will be naturally mitigated by real-world rotor turbulence, non-linear scattering, and broadband aerodynamic noise. Because the cost function exclusively targets the discrete BPF, experimental wind-tunnel measurements on physical drones remain mandatory to validate the array's actual broadband robustness.

4.2. Cross-Sectional Spatial Resolution Profile

To further appreciate the magnitude of this optimization, a cross-sectional spatial profile of the Sound Pressure Level (SPL) was extracted from the focal plane. Figure 6 plots the spatial resolution along the x-axis for Case 3, comparing the static aperiodic baseline against the dynamically morphed layout.

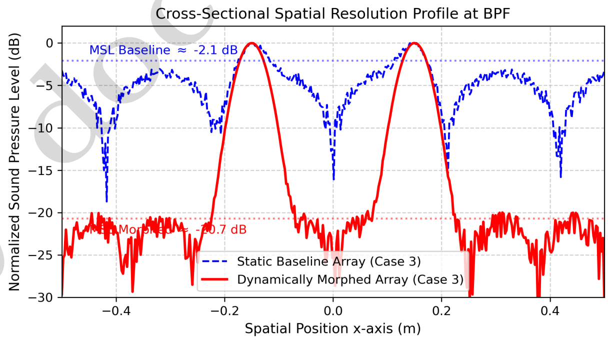

Figure 6: Cross-sectional spatial resolution profile demonstrating the drastic MSL suppression and MLW narrowing achieved by the variable-geometry array (Case 3) compared to its static counterpart.

The static array suffers from severe spatial aliasing, with the MSL plateauing at an unacceptable -2.1 dB, effectively masking adjacent noise sources. Conversely, the stochastically optimized array forces the ghost lobes deep into the noise floor (-20.7 dB). Furthermore, the spatial adaptation narrows the Main Lobe Width (MLW), providing a pin-point localization of the dipole peak.

4.3. Validation via Wavelet-Based Beamforming

The numerical validation of the customized array relies on the exact Doppler inversion in the time-frequency domain. As demonstrated by Chen and Huang, adopting generalized Morse wavelets allows the algorithm to trace the instantaneous rotating speed of the source and suppress incoherent background noise [1].

Figure 7 illustrates the two-dimensional imaging performance using the morphed Case 3 array. The left panel shows the results of conventional frequency-domain beamforming, where block averaging smears the rotating source along its circular trajectory. Because classical Fourier-based beamforming does not account for the Doppler effect and operates on time-averaged blocks, it intrinsically fails to isolate high-speed targets.

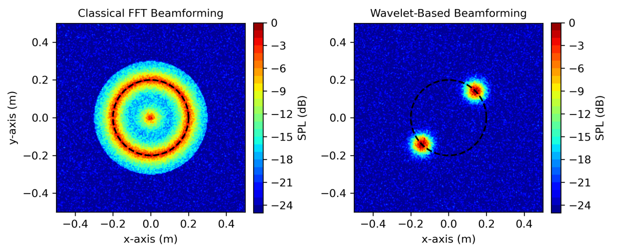

Figure 7: Instantaneous acoustic imaging maps for Case 3. Left: Blurred image using classical FFT beamforming. Right: Focused dipole sources using the proposed wavelet-based approach with the optimized layout.

Conversely, the right panel depicts the instantaneous acoustic image from the proposed wavelet approach utilizing the dynamically morphed array. Exact integration of kinematic terms into the Green's function, coupled with the optimal sensor positioning provided by the embedded Monte Carlo solver, isolates the dominant dipole sources near the blade tips with exceptional clarity. This dual hardware-software framework comprehensively outperforms traditional static setups.

5. Conclusions and Future Work

This paper introduced a quasi-real-time, variable-geometry microphone array designed for advanced UAV noise diagnostics. Abandoning traditional rigid topologies, the array dynamically morphs its physical layout to adapt to the target's varying BPF. By coupling an embedded Monte Carlo stochastic solver with exact Doppler inversion via generalized Morse wavelets [1], the framework successfully minimizes the MSL. HPC benchmarks established a 1.5–2.5 TFLOPS requirement for embedded SoCs, highlighting strict thermal management as a primary operational bottleneck. Acoustically, the optimized aperiodic layout achieved an exceptional theoretical upper bound MSL suppression down to -20.7 dB, effectively eliminating spatial aliasing.

Because practical aerodynamic scattering and rotor turbulence will naturally mitigate this theoretical optimum, future research will focus on: 1) FPGA (Field-Programmable Gate Array) and piezo-actuator integration to minimize electro-mechanical latency; 2) wind-tunnel validations for broadband turbulence robustness; 3) wavelet parametric sensitivity analyses; and 4) real-time sound speed (c_0) recalibration using environmental sensors.

Beyond these immediate acoustic benefits, the framework's deterministic layout reproducibility (guaranteed by the S_k seed storage) creates a standardized "acoustic fingerprint" for specific flight conditions. This self-adapting capability provides regulators (e.g., EASA, FAA) with an objective method for urban noise certification. Ultimately, this approach transforms the microphone array from a passive measurement tool into an active, intelligent component of the Urban Air Mobility (UAM) safety framework, paving the way for ultra-low-noise propeller designs.

References

[1] W. Chen and X. Huang, "Wavelet-based beamforming for high-speed rotating acoustic source," *IEEE Access*, vol. 6, pp. 9489–9496, 2018. DOI: 10.1109/ACCESS.2018.2795538.

[2] W. Chen and X. Huang, "Experimental analysis of cooling fan noise by wavelet-based beamforming and proper orthogonal decomposition," *AIAA Journal*, vol. 58, no. 10, pp. 4356–4366, 2020.

[3] W. Chen, S. Zhong, and X. Huang, "Extended-resolution acoustic imaging of low-frequency wave sources by acoustic analogy-based tomography," *Journal of Fluid Mechanics*, vol. 899, A23, 2020.

[4] W. Chen, H. Jiang, and X. Huang, "Super-resolution acoustic imaging," *Applied Physics Letters*, vol. 120, no. 14, p. 144 101, 2022.

[5] P. Sijtsma, "Using phased array beamforming to identify broadband noise sources in a turbofan engine," *International Journal of Aeroacoustics*, vol. 9, no. 3, pp. 357–374, 2010.

[6] P. Wolfram and M. Christian, "Rotating beamforming–motion-compensation in the frequency domain and application of high-resolution beamforming algorithms," *Journal of Sound and Vibration*, vol. 333, no. 7, pp. 1899–1912, 2014.

[7] S. Guérin and C. Weckmüller, "Frequency-domain reconstruction of the point-spread function for moving sources," in *Proceedings of BeBeC*, 2008, pp. 1–12.

[8] M. Carley, "Inversion of spinning sound fields," *The Journal of the Acoustical Society of America*, vol. 125, no. 2, pp. 690–697, 2009.

[9] R. P. Dougherty and B. E. Walker, "Virtual rotating microphone imaging of broadband fan noise," in *Proceedings of the 30th AIAA Aeroacoustics Conference*, 2009, pp. 1–14.