

ASTRAEUS Radio Interferometry Engine: An Evidence-Aware Architecture, Prototype Evaluation, and Integration Roadmap

Aaron Singh

Abstract—This paper presents ASTRAEUS Radio Interferometry Engine, a project in a twenty-project AI infrastructure portfolio. The current repository is an active prototype with implemented behavior, automated correctness tests, and explicit boundaries around unverified hardware or production claims. We describe the system model, current implementation, goals, and evaluation method, then propose five integrations that advance the project toward reproducible system-level validation. A local audit executed 3 project tests successfully. No accelerator, deployment, or performance conclusion is inferred unless a corresponding committed benchmark or profiler artifact exists.

Index Terms—ASTRAEUS, AI infrastructure, reproducibility, prototype validation, systems evaluation, hardware–software co-design

I. INTRODUCTION

ASTRAEUS Radio Interferometry Engine (P17) addresses a bounded problem within the portfolio’s track-05-hardware-software-codesign track. Its declared status is *Active Prototype* [1]. The project validates its central mechanism before adding device-specific acceleration, production orchestration, or real-world hardware. This ordering matters because optimization without a trusted reference can make incorrect behavior appear successful.

The project goals are to establish deterministic behavior, encode correctness as tests, create machine-readable evidence, identify limiting resources through measurement, and integrate results with adjacent portfolio systems without losing provenance. The repository contains one paper named exactly after its singular folder: `P17-astraeus-radio-interferometry-engine.tex`.

II. TECHNICAL CONTEXT

Hardware–software co-design links analytical models, workload evidence, accelerator mappings, and operator tools. Modeled, simulated, and measured values must remain distinguishable so design tradeoffs are not overstated. The portfolio benchmark standard requires environment, workload, method, metric, artifact, reproduction, and limitation fields; unknown values remain explicitly unmeasured [3].

III. SYSTEM MODEL AND ARCHITECTURE

The prototype is organized around the following domain model:

$$X[k] = \sum_{n=0}^{N-1} x[n]e^{-j2\pi kn/N}, \quad \hat{k} = \arg \max_k |X[k]| \quad (1)$$

The equation is a design and test abstraction rather than a claimed empirical law. It supports invariants and expected-value checks while later implementations replace synthetic inputs with representative workloads or devices.

The software architecture contains an input/configuration layer, a deterministic core, validation and evidence output, and a local visualization. The principal inspected source artifacts are `python/radio_reference.py`. Unsupported real-world conditions are surfaced as limitations rather than silently simulated.

IV. DETAILED SCRIPT OPERATION AND RATIONALE

The radio reference generates a known sinusoidal tone, computes a NumPy FFT, identifies the strongest frequency bin, and verifies expected length and input validity. It provides a CPU correctness oracle before adding antenna correlation, recorded radio data, oneMKL, SYCL, or XPU execution.

The execution path is:

- 1) Validate sample count, sample rate, and tone frequency.
- 2) Generate a deterministic sinusoidal input.
- 3) Compute the CPU FFT and magnitude spectrum.
- 4) Find the maximum-magnitude frequency bin.
- 5) Compare the detected bin with the known injected tone.

V. IMPLEMENTED PROTOTYPE

The metadata and source audit found these completed features [2]:

- Synthetic tone generator
- NumPy CPU FFT reference
- Peak-bin correctness tests

`python3 -m unittest discover -s tests -p 'test_*.py'` completed successfully with 3 tests. The inspected test artifacts are `tests/test_radio.py`. This is evidence of local correctness for encoded cases, not production scale or hardware performance.

VI. TESTING METHODOLOGY AND OBSERVED RESULTS

Testing uses Python’s standard `unittest` discovery and exercises the public behavior of the reference implementation. The audit reran the suite from the project folder with `python3 -m unittest discover -s tests -p 'test_*.py'`. All 3 discovered tests passed. The result establishes correctness only for the encoded

TABLE I
IMPLEMENTATION ARTIFACTS AND WHY THEY EXIST

Artifact	Observed responsibility	Engineering rationale
python/radio_reference.py	tone, strongest_bin, main	Implements the inspectable, unit-tested project core.
scripts/reproduce.sh	Fixed test and demonstration entry point	Gives another developer one command for local reproduction.
PROJECT.yaml	and Status, completed work, planned work, and claim boundaries Local evidence and status visualization	Separates declared intent from evidence-backed implementation.
ANALYSIS.md		Makes outputs inspectable without upgrading simulation into a hardware claim.
streamlit_app.py		

TABLE II
AUDITED TEST MATRIX

Test artifact and case	Behavior being checked	Observed result
tests/test_radio.py:test_peak	Peak.	Pass (local)
tests/test_radio.py:test_length	Length.	Pass (local)
tests/test_radio.py:test_invalid	Invalid.	Pass (local)

local cases; it does not establish accelerator correctness, real-device behavior, production reliability, or benchmark completion.

No numerical performance result is promoted by this test run. Where scripts emit JSON or JSONL, those outputs remain raw or simulation-specific until a reviewed summary includes hardware, software, workload, warm-up, repetition, correctness threshold, Git revision, and limitations.

VII. CLAIM BOUNDARIES AND RISKS

The project records these unverified or excluded claims:

- No interferometry, radio dataset, oneMKL, or SYCL result

The main risk is confusing synthetic or modeled behavior with deployed-system behavior. Other risks include incomplete workloads, platform-dependent timing, missing failure injection, and interfaces not yet exercised across device boundaries. Performance claims require a reviewed record meeting the portfolio standard.

VIII. EVALUATION PLAN

Evaluation proceeds through correctness tests, deterministic reproduction with Git and environment metadata, repeated benchmarks reporting latency/throughput/memory/error metrics, and a named profiler capture tied to exact hardware and source revision. Success requires reference equivalence within a documented tolerance, preservation of safety and resource invariants, clear failures, and evidence reproducible from a clean environment.

IX. GOALS, MILESTONES, AND SUCCESS CRITERIA

The project goals are staged so that correctness precedes performance and integration. A goal is complete only when its proof artifact is committed or otherwise reviewable; prose or a simulated number alone is insufficient.

X. FUTURE WORK AND INTEGRATIONS

The five project-specific next steps are:

- 1) Two-antenna correlation
- 2) Recorded radio sample
- 3) oneMKL/SYCL implementation
- 4) Add calibration, windowing, synchronization, and uncertainty estimates.
- 5) Publish CPU/XPU evidence to P19 after SYCL validation.

The early items complete declared evidence; the later items connect downstream portfolio consumers. Each integration should add tests and a reviewable artifact such as JSONL evidence, a report, profiler capture, deployment manifest, trace, or labeled data set.

XI. CONCLUSION

ASTRAEUS Radio Interferometry Engine is an evidence-aware active prototype: its implemented behavior and tests are real, while unbuilt hardware, deployment, and performance goals remain labeled. Completing the five integrations in dependency order will advance it from a learning artifact toward a credible portfolio component.

REFERENCES

- [1] Aaron Singh, “ASTRAEUS Radio Interferometry Engine PROJECT.yaml,” local portfolio repository.
- [2] Aaron Singh, “ASTRAEUS Radio Interferometry Engine: README, ANALYSIS, source, and test artifacts,” local portfolio repository.
- [3] Aaron Singh, “AI Infrastructure Portfolio Benchmark Standard,” local portfolio repository.

TABLE III
DECLARED STATUS VERSUS AUDITED EVIDENCE

Audit field	Finding	Evidence source
Declared status	Active Prototype	PROJECT.yaml
Evidence-backed status	Active local prototype; 3 tests passed	Source plus local unittest run
Accepted measured results	None recorded in measured_results	PROJECT.yaml
Mismatch / claim boundary	No interferometry, radio dataset, oneMKL, or SYCL result	PROJECT.yaml and ANALYSIS.md
Next proof required	Two-antenna correlation; Recorded radio sample	Planned features

TABLE IV
PROJECT GOALS AND REQUIRED PROOF

ID	Goal	Completion evidence
G1	Two-antenna correlation	Passing tests and a reviewed source artifact
G2	Recorded radio sample	Machine-readable result with reproduction metadata
G3	oneMKL/SYCL implementation	Reference-equivalence or domain-correctness report
G4	Implement two-antenna correlation	Named profiler, deployment, or integration artifact
G5	Validate oneMKL/SYCL against CPU results	Dashboard/report link preserving provenance and limitations