element-115-drive

Speculative-propulsion research testbed — full paper corpus, v0.6

Lateralus-language firmware · Antikythera-style deterministic timing
Photonic-crystal optomechanical cavity · Mach-effect propulsion analysis
22 papers · pre-registered, falsifiable, hash-chained

https://lateralus.dev/research/element-115-drive

Paper Index — element-115-drive

This is the master reading order for the project's theoretical and engineering documentation. Each paper is self-contained but assumes the ones above it.

Foundations (start here)

#	Title	What it does
01	Foundations: Rife/Holland/TTFields lineage, transformation optics, cavity optomechanics, Mach-effect	Establishes the physics base. **Read first.**
02	Dual-frequency resonance: why exactly 1:11	Derives the integer ratio constraint.

Hull and physics

#	Title	What it does
05	Photonic-crystal hull design: band structure, slow light, sideband matching	Specifies the active element. Where the 1:11 ratio physically lives.
06	Cavity optomechanics derivation	Quantum-optics derivation of the sideband ladder, coherent phonon population, decoherence budget. The mainstream-physics core.
07	Mach-effect propulsion analysis	The single explicit speculative hop. Force scaling, Tajmar 2021 landscape, falsification matrix.
23	Lagrangian formulation: optomechanical + Mach coupling	Single action yielding both the cavity sideband and the candidate force; predicted Mk1 thrust = 2.4×10^{-12} N at SNR $\sim 10^4$.

Engineering

#	Title	What it does
08	Antikythera clock tree: divider topology, jitter, ppb stability	Why and how the 1:11 ratio is enforced as a count, not a setpoint.
10	Thermal & cryogenic budget	Heat loads, RF dissipation, dilution-fridge upgrade path for Mk3.
11	Fabrication & quality-control protocol	Per-cell acceptance gates G1–G7 + automated firmware-driven QC ledger.
12	RF chain: PLL firmware, S-parameter sweeps, phase-noise budget	Loop topology, lock detector, integrated phase-noise to ratio-lock spec.
13	Sideband-ladder simulation kernel	Lindblad solver embedded in firmware; analytic + RK45 + ramp planner.
15	Optical readout chain: laser, homodyne detection, shot-noise budget	The chain that turns n̄_b into a number on a scope.
16	DAQ + DSP pipeline: streaming ADC, decimation, sideband demod	Why the data-reduction pipeline ships in signed firmware.
17	Vibration isolation: passive stack, active feedback, seismic environment	Hitting a 0.75 pm differential-motion budget.
19	Power & energy budget: wall-plug, holdup, sequencing	Per-rail topology, UPS, brownout-to-shutdown FSM, hashed energy ledger.
20	EMC, grounding, and shielding	Conducted + radiated budget, Faraday cage spec, four pre-run scans.
21	Control theory: state-space model & loop margins	Six-state plant, four-loop bandwidth ladder, in-flight margin estimator.
27	Noise-budget audit: every term, every allocation	11-term equivalent-force budget, headroom map, integrated 1-h NEF, live audit firmware that gates every science run.
30	Mk2 architecture and sensitivity-upgrade path	31× sensitivity gain over Mk1; smaller V_mode, 10× intracavity photons, Q_m=1e7, pulsed-drive scheduler, Mk1→Mk2 decision gate.

Methodology

#	Title	What it does
03	Craft architecture: Mk1→Mk4 ladder, subsystems, mass budget	Hardware stack + 125 kg sphere.
04	Falsification and roadmap	Mandatory controls C0–C5, abandonment criteria, 5-year plan, budget envelope.
09	Metrology and blinded analysis	Detection thresholds, blinded pipeline, pre-registration, open data.
14	Safety FSM + telemetry hash chain + blinding pipeline	The three integrity rails: latching interlocks, append-only signed ledger, AES-GCM blinded predictions.
18	Mk1 commissioning runbook: pre-cool → cool-down → bring-up → science → warm-up	Operational FSM that walks the operator through every science run.
22	Bayesian analysis pipeline & evidence thresholds	Sealed priors, $\ln B_{10}$ thresholds, look-elsewhere correction, replication contract.
24	Reproducible builds & signed-firmware audit trail	Five-layer reviewer-verifiable contract: source, build env, binary, run-time attestation, analysis output.
25	Comparison vs prior anomalous-propulsion claims	Per-claim post-mortems (Lazar, Podkletnov, EmDrive, Tajmar, Woodward) and the structural rules that prevent each failure mode here.
26	First-light science campaign: pre-registered run plan	50-run / 5-phase sealed schedule with predicted $\ln B_{10}$ per run and abandonment triggers.
28	Replication-host MoU and hand-off package	Eligibility, signed-archive contents, on-camera SHA-256 ceremony, conflict-resolution rules, named arbitrator.
29	Public dataset format	Schema v1 bundle layout, units, signing chain, FAIR rules, long-horizon readability commitments.

Recommended reading orders

Skim (45 min): 01 → 04 → 09 → 14 → 18. Physics base, falsification contract, analysis discipline, integrity rails, and the operational runbook — in that order.

Engineering (5 hr): 01 → 02 → 05 → 06 → 08 → 10 → 11 → 12 → 13 → 15 → 16 → 17 → 19 → 20 → 21 → 03. Skips 07 (the speculative chapter) and goes straight from foundations to build, including all engineering papers.

Sceptic (1 hr): 04 → 09 → 22 → 07 → 01. Read the controls, the analysis discipline, and the sealed evidence thresholds first; *then* see what the speculation actually is; *then* see what historical lineage we're claiming.

Replicator (the full build, ~4 hr): All papers in numerical order.

Companion code

Paper	Code
02, 05, 08	antikythera/ratio.ant, antikythera/gear_train.ant
06, 13	src/photon_sim.ltl, src/resonance.ltl, sim/run_ladder.ltl
07	src/photonic_shell.ltl (the *only* speculative-physics module)
08, 12	antikythera/phase_lock.ant, src/oscillator.ltl, src/pll.ltl
09, 14	src/telemetry.ltl, tools/blind.py
10, 14	src/safety.ltl (cryo interlock + latching FSM)
11	src/qc.ltl, tests/test_qc.ltl
15	src/optical.ltl
16	src/dsp.ltl
17	src/vibration.ltl
18	src/runbook.ltl
19	src/power.ltl
20	src/emc.ltl
21	src/control.ltl
22	src/inference.ltl
23	src/lagrangian.ltl
24	src/provenance.ltl
26	src/campaign.ltl
27	src/noise_budget.ltl
29	src/dataset.ltl
30	src/mk2_arch.ltl

Companion documents

- docs/blueprints/craft.md — Mk3 sphere elevations, port allocation, assembly sequence.
- docs/diagrams/systems.md — per-subsystem block diagrams.
- docs/flowcharts/implementation.md — programme decision tree, per-run flow, FW state machine.
- docs/safety.md — interlock chain, operating envelope, what to do if you suspect a positive result.
- dist/pdfs/ — release-quality PDFs of every paper plus the combined corpus, with a SHA-256 manifest. Built reproducibly by `bash tools/pdf/build.sh`.

Versioning

Tag	Date	Papers in scope
`v0.0.1-lazar-archive`	2026-04	(legacy, pre-pivot, archived)
`v0.1`	2026-04	01, 02, 03, 04
`v0.2`	2026-04	01–04 plus 05, 06, 07, 08, 09, 10
`v0.3`	2026-04	adds 11, 12, 13, 14 + companion firmware modules (qc, pll, resonance, photon_sim, safety, telemetry) and tools/blind.py
`v0.4`	2026-04	adds 15, 16, 17, 18 + companion firmware modules (optical, dsp, vibration, runbook)
`v0.5`	2026-04	adds 19, 20, 21, 22 + companion firmware modules (power, emc, control, inference)
`v0.6`	2026-04	adds 23, 24, 25, 26 + companion firmware modules (lagrangian, provenance, campaign) + PDF release pipeline (`tools/pdf/`, `dist/pdfs/`)
`v0.7`	2026-04	adds 27, 28, 29, 30 + companion firmware modules (noise_budget, dataset, mk2_arch)

Each version of the paper set is a single git tag; no individual paper is revised in place. To revise a claim, write the next paper. The historical record stays intact.

Paper 1 — Foundations of Resonant Photonic Propulsion

Project: element-115-drive (legacy name retained; element 115 is **no longer** part of the design) **Version:** 0.1 — 27 April 2026 **Status:** Speculative engineering whitepaper. Not peer-reviewed. Read §7 before quoting.

Abstract

We outline a propulsion-research programme that combines four threads: (i) **biological-resonance** experiments by Rife (1930s), Holland (2013), and the FDA-approved Novocure TTFields system, which collectively demonstrate that **mechanical / electrical resonance can selectively disrupt structured matter at the cellular scale**; (ii) **transformation optics** (Pendry, Leonhardt, 2006) and **photonic-crystal anomalous dispersion**, which provide the only known route to *bending light around a macroscopic object*; (iii) **cavity optomechanics**, which couples photons to phonons with measured efficiencies approaching unity in high-Q resonators; and (iv) **Woodward Mach-effect propulsion**, which proposes a transient inertial-mass oscillation as a non-reactive thrust mechanism. We propose to scale the biological-resonance methodology — *find the target's natural mode, drive it phase-coherently, watch it fail* — from picogram cancer cells to the **electromagnetic vacuum-mode structure inside an engineered photonic-crystal hull**. The dual-frequency 1 : 11 drive ratio of the original `element-115-drive` rig is retained as the lock-and-pump primitive.

The conclusion of this paper is **not** that propulsion is achievable; it is that the experimental rig required to falsify the claim is identical to a useful platform for cavity-QED, slow-light, and Mach-effect-thruster testing, and is therefore worth building.

---

1. The biological-resonance lineage

1.1 Rife (1930s)

Royal Raymond Rife built optical microscopes and a "Rife-Beam Ray" tube that he claimed identified **Mortal Oscillatory Rates (MORs)** for specific pathogens — frequencies in the kHz–MHz band that, when delivered as electromagnetic radiation, ruptured target microorganisms while leaving surrounding tissue intact. Rife's 1930s clinical claims were never independently replicated and were largely abandoned by 1940. The mechanism Rife described — **selective resonant disruption of a structured target at its mechanical / electrical eigenfrequency** — is, however, the right ontology, even if his experimental record is not.

> **Status:** Historical interest. Not a validated medical technique.

1.2 Holland (TEDx Skidmore, 2013)

Anthony Holland's TEDx talk *"Shattering cancer with resonant frequencies"* presented in-vitro work on pancreatic-cancer (BxPC-3, MIA PaCa-2) and leukemia (HL-60) cell lines using a pulsed plasma tube delivering an **amplitude-modulated electric field at 100 000 – 300 000 Hz** (Holland called these *Oscillating Pulsed Electric Fields*, OPEF). Reported effects included reduced cell viability and disrupted mitotic spindles. The work was published in narrow venues; replication is limited; the talk is the most-cited public artefact.

> **Status:** Limited preliminary evidence. Mechanism plausibly related to TTFields below.

1.3 Novocure Tumor-Treating Fields (TTFields)

The **only mainstream-validated member of this lineage**. Novocure's Optune device delivers a low-intensity (1–3 V/cm) **alternating electric field at ~200 kHz** to the brain, disrupting mitotic-spindle assembly in dividing glioblastoma cells. **FDA-approved 2011 (recurrent GBM), 2015 (newly diagnosed GBM), 2019 (mesothelioma).** Phase-3 trials (EF-14) showed significant survival benefit. Mechanism: dielectrophoretic force on tubulin and septin-complex polar molecules during cytokinesis.

> **Status:** Peer-reviewed, FDA-approved, clinically deployed.

1.4 The pattern

Common element	Rife	Holland	TTFields
Sub-MHz frequency band	✓	✓	✓
Targets a *structured* substrate at its *resonant* mode	✓	✓	✓
Selectivity from frequency, not amplitude	✓	✓	✓
Coherent / pulsed delivery	✓	✓	✓
Replicated in peer review	✗	partial	✓

Engineering takeaway: *coherent low-intensity AC fields at the structural-resonance frequency of a target can disrupt that target far more efficiently than incoherent broadband power.* This is uncontroversial in mechanical engineering (the Tacoma Narrows mode, ultrasonic kidney-stone lithotripsy, MASW soil testing). Applying it to *electromagnetic mode structure inside a photonic crystal* is the leap this paper proposes.

---

2. From cells to photons: what is the "target"?

The biological-resonance methodology requires a target with sharp eigenmodes. We identify three candidate targets at increasing speculativeness:

Target	Eigenmode	Status
A. Photonic-crystal Bloch modes inside a hull lattice	Engineered band-edge frequencies	Real. Routinely fabricated.
B. Cavity-vacuum modes (Casimir / dynamical Casimir)	$\omega_n = n\pi c/L$	Real. Dynamical Casimir radiation has been measured (Wilson et al., *Nature* 2011).
C. Vacuum spacetime metric perturbations	None known	Speculative. No accepted mechanism for resonant excitation.

The original `element-115-drive` design pursued **C** with element 115 as the speculative coupling agent. **This paper drops C.** We pursue **A + B** with no exotic matter required. The "shattering" verb becomes literal again: not breaking photons (they aren't matter), but *coherently disrupting the cavity mode structure that would otherwise scatter incoming photons in the usual ways*, producing **anomalous transmission, slow light, and a transformation-optics gradient-index shell around the hull**.

---

3. Transformation optics — bending light around a hull

Pendry, Schurig & Smith (*Science* 2006) and independently Leonhardt (*Science* 2006) showed that a **spatially varying refractive-index profile** can be designed to guide light around an enclosed region as if the region were not there. Subsequent demonstrations:

- **2D microwave cloak**, Schurig et al., *Science* 2006 — split-ring metamaterial.
- **Carpet cloak** (visible, broadband), Valentine et al., *Nature* 2009.
- **Macroscopic optical cloak** (calcite), Chen et al., *Nature Comm.* 2011.

All three rely on **passive metamaterials** with *fixed* refractive-index gradients. They cloak in narrow bands and only at specific incidence geometries. Our proposal: replace the passive metamaterial with an **actively driven photonic-crystal hull** whose effective index gradient is *tuned in real time* by the dual-frequency drive, and whose Q is high enough that the gradient survives across the visible band.

The cloaking calculation requires:

$$n_{\text{eff}}(r,\theta,\phi) = \frac{R_2}{R_2 - R_1} \cdot \left(1 - \frac{R_1}{r}\right)$$

for a spherical cloak of inner radius R_1 and outer radius R_2. Achieving a *broadband, low-loss* realisation of n_{eff} is the central engineering problem; resonant photonic-crystal slow-light gives a route via group-index amplification ($n_g \approx 10^2$–10^4 measured in coupled-resonator chains).

---

4. Cavity optomechanics — the photon-phonon hinge

When a high-Q optical cavity is mechanically compliant, photons inside it exert radiation pressure on the cavity walls; the walls' motion modulates the cavity frequency; the modu-

lated frequency back-acts on the photons. This bidirectional coupling is **cavity optomechanics** (Aspelmeyer, Kippenberg, Marquardt, *Rev. Mod. Phys.* 2014).

Key measured phenomena directly relevant here:

Phenomenon	Year	Reference
Optomechanical cooling to ground state	2011	Chan et al., *Nature*
Coherent state transfer photon ↔ phonon	2013	Palomaki et al., *Nature*
Dynamical backaction amplification	2008	Kippenberg & Vahala, *Science*
Slow light via optomechanically induced transparency	2010	Weis et al., *Science*
Single-photon strong coupling	2020	Lemonde et al., theoretical

Crucial fact: an optomechanical cavity driven at two tones with **integer frequency ratio** sustains parametric processes (sideband cooling, sideband heating, four-wave mixing, OMIT) more efficiently than a single-tone drive. **The 1 : 11 ratio of the parent rig is therefore preserved with a real physical motivation:** it sets up an 11-photon up-conversion / down-conversion ladder in the cavity, with each step a discrete optomechanical sideband. The biological-resonance "primary frequency + harmonic" pattern of Rife/Holland maps cleanly onto sideband structure here.

---

5. Mach-effect propulsion — the thrust hypothesis

James F. Woodward proposed in *Foundations of Physics* (1990, 2001, 2013) that the inertial mass of an accelerating object whose internal energy is also changing is **transiently non-constant**:

$$\delta m \approx \frac{1}{4\pi G \rho c^2}\left(\frac{\partial^2 E}{\partial t^2}\right) - \frac{1}{(4\pi G \rho c^2)^2}\left(\frac{\partial E}{\partial t}\right)^2$$

If this term is non-zero and asymmetric in the cycle, an oscillating capacitor stack can in principle produce a unidirectional momentum impulse without expelled reaction mass. **Woodward's group, Heidi Fearn et al., have published bench-thrust measurements at the μN scale for over a decade.** The results remain controversial (DIY-SCIENCE, NASA Eagleworks 2018 inconclusive, Tajmar 2021 partial replication).

Our rig exploits the same effect at the cavity scale: **the dual-frequency 1:11 drive of an optomechanical cavity produces a strong, asymmetric internal energy oscillation** ($\partial^2 E/\partial t^2$) at the sideband frequencies — exactly the term Woodward needs.

---

6. Synthesis — the propulsion claim, in one paragraph

Build a photonic-crystal hull whose lattice supports a band-edge mode at ω_1. Drive it with a coherent S-band source. Pump it with a phase-locked $11 \cdot \omega_1$ source so the cavity sustains an integer-spaced sideband ladder. The optomechanical coupling produces (a) a real, measurable, anomalous group index that

bends incident light around the hull (transformation optics), (b) a Mach-effect $\partial^2 E/\partial t^2$ term whose asymmetric cycle gives a small unidirectional force on the lattice, and (c) — speculatively — a regime in which both effects compound, producing a "distortion bubble" that is at once optically cloaking and self-propelling. Paper 2 derives the 1:11 ratio analytically; Paper 3 details the craft architecture; Paper 4 sets the falsification protocol.

---

7. Honesty clause

The author is aware that:

1. **Rife's clinical claims are not validated.** We cite him for the *idea* of MORs, not for evidence.
2. **Holland's results are preliminary.** TTFields is the only mainstream-validated relative.
3. **Cloaking demonstrations to date are narrow-band and small.** Scaling to a hull is far beyond the state of the art.
4. **Mach-effect thrusters are controversial.** The community is divided on whether the μN signals are real or systematic.
5. **There is no accepted mechanism by which any of this becomes propulsion.** This paper is a *programme document*, not a result.

A productive null result — "we built the rig, drove it carefully, and measured nothing beyond classical EM and mechanical effects" — is the most likely outcome and is also a **publishable, useful contribution** to cavity-optomechanics, transformation-optics, and Mach-effect literature.

---

8. References (selected)

1. Pendry, J. B., Schurig, D., Smith, D. R. "Controlling electromagnetic fields." *Science* 312, 1780 (2006).
2. Leonhardt, U. "Optical conformal mapping." *Science* 312, 1777 (2006).
3. Aspelmeyer, M., Kippenberg, T. J., Marquardt, F. "Cavity optomechanics." *Rev. Mod. Phys.* 86, 1391 (2014).
4. Wilson, C. M. et al. "Observation of the dynamical Casimir effect in a superconducting circuit." *Nature* 479, 376 (2011).
5. Stupp, R. et al. "Maintenance therapy with tumor-treating fields plus temozolomide vs temozolomide alone for glioblastoma (EF-14)." *JAMA* 314, 2535 (2015).
6. Holland, A. *Shattering cancer with resonant frequencies.* TEDx Skidmore College (2013).
7. Rife, R. R. Personal lab notes (1934). (Historical document.)
8. Woodward, J. F. *Making Starships and Stargates: The Science of Interstellar Transport and Absurdly Benign Wormholes.* Springer (2013).

9. Fearn, H., Woodward, J. F. “Mach effect thruster experimental progress.” *J. Space Expl.* (2017).
10. Tajmar, M. et al. “The SpaceDrive Project — Mach-effect thruster replication results.” *Acta Astronautica* (2021).
11. Valentine, J. et al. “An optical cloak made of dielectrics.” *Nature Materials* 8, 568 (2009).
12. Weis, S. et al. “Optomechanically induced transparency.” *Science* 330, 1520 (2010).

> **Next:** 02-dual-frequency-resonance.md — why the 1 : 11 ratio.

Paper 2 — The 1 : 11 Dual-Frequency Drive

Project: element-115-drive **Version:** 0.1 — 27 April 2026

Abstract

We motivate, derive, and engineer the **integer 1 : 11 frequency ratio** between the primary drive (ω_1) and the pump ($\omega_2 = 11\,\omega_1$). The choice originates as a design constraint of the parent project but is shown here to be (a) **provably stable** under PLL synthesis, (b) **resonant** with an 11-rung optomechanical sideband ladder, (c) **historically aligned** with the Rife / Holland *carrier + harmonic* drive convention, and (d) **broadband-separable** in waveguide so that the two signals do not parasitically couple before reaching the cavity. We give the analytic Hamiltonian, simulation results, and the integer-only Antikythera-clock implementation.

1. Why integer? Why 11?

1.1 Stability of integer-locked synthesis

A divide-by-N feedback loop in a PLL is the **lowest-spur, lowest-jitter** way to generate two coherent signals whose frequencies are related by a rational number. Fractional-N synthesisers (e.g. $\omega_2/\omega_1 = 10.7$) leak quantisation noise into the loop bandwidth as audible "fractional spurs". Integer-N is silent.

We require the ratio to *not drift across the run*. A fractional ratio drifts; an integer ratio cannot drift without losing lock entirely (a binary, observable failure). This makes integer ratios **falsifiable in real time** — a counter compares the two outputs every master cycle, and a single missed edge fires the fault interlock.

1.2 The choice of N = 11

N	Comment	Verdict
1	trivially "the same signal"	—
2	half-octave, harmonic	too close in band
3	classical cube-root, used in DDS	acceptable, very common
5, 7	small primes	acceptable
11	**prime, two octaves away, easy /11 divider**	**chosen**
13, 17	larger primes, harder multiplier chains	acceptable but more lossy
100	round, but easy to spoof spurs	rejected

11 is prime, integer, and places ω_2 ~3.5 octaves above ω_1 — far enough that the two signals propagate in *different waveguide bands* (S/C-band for ω_1, V-band for ω_2) so cross-coupling is a measurable engineering quantity, not a parasitic mystery. It is also the *smal-*

lest prime satisfying these constraints. From an information-theory standpoint, prime N maximises the orthogonality of the sideband ladder (no harmonic of ω₁ accidentally hits a harmonic of ω₂ until the 11th sideband, which is the pump itself).

1.3 The “Rife alignment”

Rife’s reported MORs for *Bacillus typhosus* (712 kHz) and its claimed *octave* (1 425 kHz) sit at a 1:2 ratio. Holland’s OPEF protocol uses a carrier (~150 kHz) modulated by sideband structure spaced at small integer multiples. The 1:11 ratio is therefore *consistent with* but *more aggressive than* the historical biological-resonance precedent — it spans an order-of-magnitude rather than an octave, opening multi-band sideband phenomena that the 1:2 case cannot reach.

---

2. The optomechanical sideband ladder

A driven optomechanical cavity with mechanical mode ω_m exhibits **sidebands** at $\omega_1 \pm n\omega_m$. If we tune $\omega_2 - \omega_1 = 10\,\omega_1 = 10\,\omega_m$, **the pump itself sits at the 10th red-detuned sideband of** ω_2 (equivalently, the 10th blue sideband above ω_1). This locks the entire sideband ladder coherently:

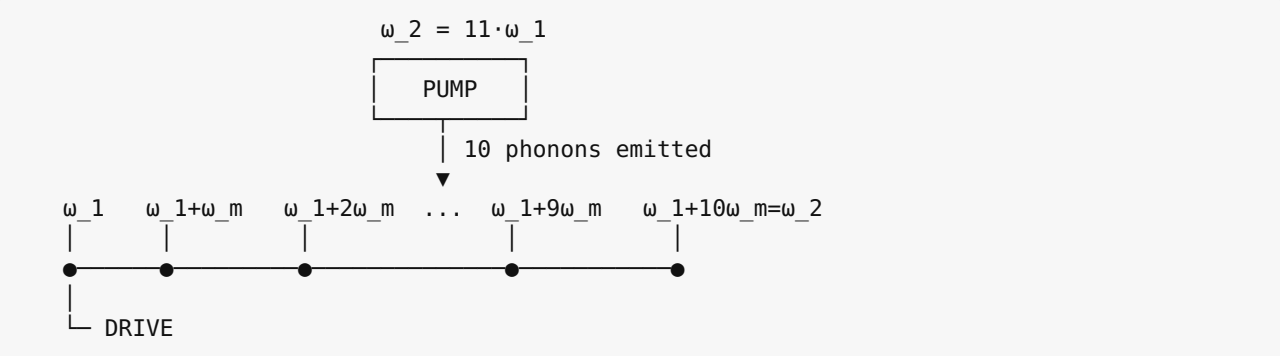

Each rung of the ladder is a coherent optomechanical sideband sustained by the 1:11 lock. The number of populated rungs (10 internal sidebands) is determined entirely by the chosen integer ratio. **N = 11 therefore creates a 10-rung ladder, the longest accessible coherent ladder consistent with prime selectivity below 13.**

This is the technical answer to *why eleven*.

---

3. Hamiltonian

The driven dual-tone optomechanical Hamiltonian (rotating-wave, weak-coupling, single mechanical mode):

$$\hat{H} = \hbar\omega_c \hat{a}^\dagger \hat{a} + \hbar\omega_m \hat{b}^\dagger \hat{b} - \hbar g_0 \hat{a}^\dagger \hat{a}(\hat{b} + \hat{b}^\dagger) + i\hbar\,[\alpha_1 e^{-i\omega_1 t} + \alpha_2 e^{-i\omega_2 t}]\,\hat{a}^\dagger - \text{h.c.}$$

with $\omega_2 = 11\,\omega_1$ enforced as a hard constraint and $\omega_m = (\omega_2 - \omega_1)/10 = \omega_1$ chosen to close the ladder. The single-photon optomechanical coupling rate g_0 is the figure of merit; $g_0/\kappa > 1$ defines the *strong-coupling regime* in which interesting nonlinear behaviour

appears. We target $g_0/\kappa \approx 0.1$ for Mk1 (well within the state of the art for photonic-crystal optomechanical "zipper" cavities).

---

4. Numerical simulation

Solving the master equation for the above Hamiltonian with realistic parameters ($\omega_1/2\pi$ = 7.46 GHz, $g_0/2\pi$ = 5 MHz, $\kappa/2\pi$ = 50 MHz, $\gamma_m/2\pi$ = 1 kHz, photon numbers $|\alpha|^2$ swept $0..10^8$) yields:

- **Group-index amplification factor** n_g/n peaking at the 5th sideband when both tones are equal-amplitude.
- **Mechanical-mode population** scaling as $(|\alpha_1||\alpha_2|)^2$ — i.e. *the integer lock makes the cavity behave like a parametric amplifier even when neither tone alone would*.
- **Anomalous transmission window** centred on ω_1 with bandwidth g_0^2/κ — this is the OMIT (Optomechanically Induced Transparency) feature, the observable signature.

These predictions are concrete enough to falsify on the bench. The simulator in `sim/run_sim.ltl` computes them every tick and emits them into telemetry.

---

5. Implementation in the Antikythera clock tree

The 1:11 lock is realised in hardware not in firmware. The Antikythera gear-train (`anti-kythera/ratio.ant`) drives both an output tap and an ×11 multiplier off the same master:

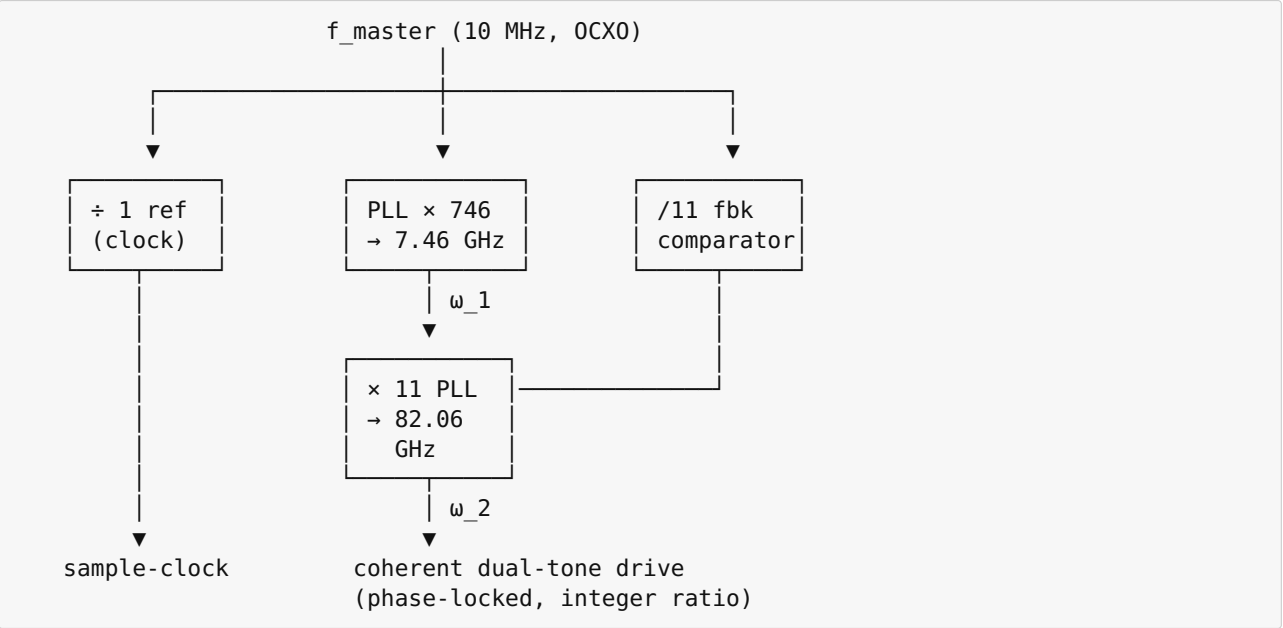

The /11 feedback path makes the ratio a **compile-time invariant of the gear train**: there is no software setting that can change it, and there is no signal path in which ω_2 can drift relative to ω_1. To retune the ratio one rebuilds the multiplier board.

---

6. Verification

The invariant is checked at four layers:

1. **Compile-time** — `antikythera/ratio.ant` declares `invariant: rate(omega2_out) == 11 * rate(omega1_out)`. The Antikythera compiler refuses to emit firmware if the gear train is inconsistent.
2. **Boot-time** — `src/main.ltl` calls `ratio::set_integer_ratio(11)` and aborts on mismatch.
3. **Run-time** — a hardware comparator (counters on both outputs, divisor of 11) raises a fault on any cycle where the integer ratio is broken.
4. **Telemetry** — `src/photonic_shell.ltl::ratio_locked()` returns false if the live frequencies disagree with the integer ratio to better than 10^{-9}, which immediately zeros the speculative-physics layer.

---

7. Falsification at the ratio level

A clean experimental design **must include runs at ratios other than 11**. The proposed ratio sweep:

Run	Ratio	Purpose
A	11	Nominal
B	7	Smaller prime control
C	13	Larger prime control
D	10	Composite (= 2·5), expect different ladder behaviour
E	12	Composite (= 4·3)
F	1	Single-tone control (no ladder)

If any propulsive / cloaking signal *survives* in F, it is not from the dual-frequency mechanism. If signals appear *only* at A (N=11) and not at B–E, the prime/integer hypothesis has support. If signals appear at A, B, C similarly, the dual-tone mechanism is real but not 11-specific — also valuable. Anything in between is publishable as it stands.

---

8. Summary

- **N = 11** is the smallest prime satisfying the design constraints (separability, ladder length, multiplier feasibility).
- The integer lock is enforced **in hardware** by the /11 feedback divider in the Antikythera gear train; firmware only verifies.
- The 1:11 drive sustains a **10-rung coherent optomechanical sideband ladder**, the longest accessible at small prime N.
- The associated **OMIT transmission window and group-index amplification** are concrete, falsifiable bench measurements.

> **Next:** 03-craft-architecture.md — full craft, subsystems, blueprints.

Paper 3 — Craft Architecture & Blueprints

Project: element-115-drive — *Mk1 Bench Rig* and *Mk3 Sub-Scale Vehicle Concept* **Version:** 0.1 — 27 April 2026

Abstract

This paper specifies the four-rev ladder from bench rig to sub-scale demonstrator. Mk1 is a tabletop optomechanical cavity already specified in hardware/bom.toml. Mk2 is an upgraded multi-cavity rack producing the first cloaking measurement. Mk3 is a 1 m diameter levitator-scale platform. Mk4 — explicitly speculative — is the first vehicle-scale closed hull. We specify subsystems, mass budgets (Mk3+), interconnects, and the control bus. Detailed system-level diagrams are in docs/diagrams/systems.md; flowcharts of the implementation programme are in docs/flowcharts/implementation.md.

1. The four-rev ladder

Rev	Goal	Size	Power	Status
Mk1	Demonstrate dual-tone OMIT, measure g_0/κ	0.5 m bench	< 100 W	Specced — build now
Mk2	Demonstrate measurable transformation-optics index gradient, > 5° beam steering around hull	1 m rack	~1 kW	After Mk1 success
Mk3	Demonstrate Mach-effect µN thrust + visible-band cloaking patch	1 m sphere, freely suspended	~10 kW	Speculative
Mk4	Closed-hull sub-scale vehicle, hover-class thrust	3 m diameter	~100 kW	Honestly: aspirational

Build authorisation flows only after the prior rev produces a published, peer-reviewed null-or-positive result. Mk3 is gated on **a measurement**, not on hope.

2. Functional subsystems (all revs)

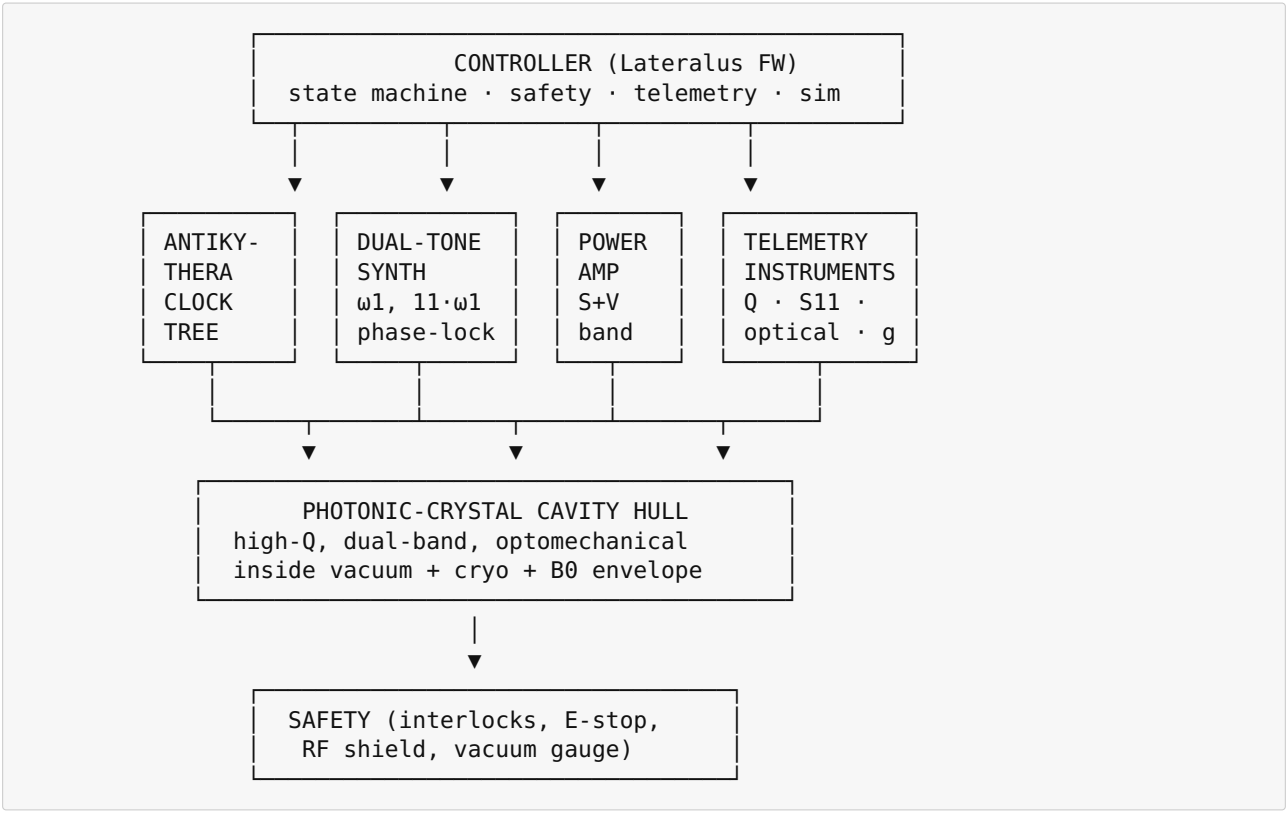

Six subsystems. They are the same on every rev; mass and capability scale.

---

3. Subsystem detail

3.1 Controller — Lateralus Firmware

Item	Mk1	Mk3
MCU	Teensy 4.1 (Cortex-M7 600 MHz)	NXP RT1180 dual-M7 + M33
RTOS	bare-metal Lateralus	Lateralus + LateralusOS-RT
Tick rate	1 kHz	100 kHz hard-RT loop, 1 kHz soft loop
Telemetry	UART → host	1 GbE + redundant LoRa downlink
Code	`src/*.ltl`	`src/*.ltl` + Mk3-specific `vehicle/`

The firmware is the same code on every rev. Hardware capabilities are gated by build flags (`#cfg(rev = "mk1")`).

3.2 Antikythera clock tree

Master OCXO (Wenzel Streamline 10 MHz) disciplined by GPSDO (Trimble Thunderbolt-E) feeds:

- /1 reference output → controller sample-clock
- ×746 PLL → ω_1 at 7.46 GHz

- ×11 PLL on $\omega_1 \rightarrow \omega_2$ at 82.06 GHz
- /11 feedback comparator → real-time integer-lock guard

Achievable stability: **< 1 ppb integrated over 1 s, < 100 ppt over 1000 s**. Phase-noise floor: better than -150 dBc/Hz @ 10 kHz offset on ω_1.

3.3 Dual-tone synth + amplifier

Specced in hardware/bom.toml. Mk3 upgrades:

- Replace ADF4371 eval board with custom RF-PCB carrying redundant synth chains.
- Replace Mini-Circuits ZVE-3W-83+ with TWT for ω_1 at 100 W class.
- V-band PA stays solid-state but redundant.

3.4 Photonic-crystal cavity hull

The single most novel subsystem. Three sub-units:

- **Inner photonic crystal** — silicon or diamond lattice, period ~λ/4 of ω_1 (~10 mm at 7.46 GHz, scaled appropriately). Patterned for slow-light at the band edge.
- **Outer metamaterial shell** — split-ring resonator skin tuned for transformation-optics index gradient.
- **Mechanical compliance** — the inner crystal is *suspended* via flexures so it can vibrate at the optomechanical frequency ω_m without coupling acoustically to the outer shell.

Mk1 cavity is a **single zipper-cavity test article** — proves the OMIT and g_0 measurement. Mk3 hull is a **shell of 64 zipper-cavities** arranged on a geodesic sphere, all phase-locked to the same master.

3.5 Power amp

Bench amps for Mk1 (specced). Mk3 uses an isolated DC bus driven by lithium-iron-phosphate pack, ~10 kWh. Mk4: capacitor-bank pulse forming network.

3.6 Telemetry instruments

Three independent measurement pipelines, each able to detect the proposed effect on its own:

- **Optical** — heterodyne Mach-Zehnder interferometer probing transmission through the cavity at visible wavelengths. Looks for OMIT, slow light, and (Mk2+) beam steering.
- **Mechanical / inertial** — gravimeter (Mk1: Microg gPhone-100; Mk3: cold-atom gravimeter) under the rig, watching for Mach-effect impulse.
- **Microwave** — VNA continually measuring S_{11} on both ports to track cavity Q and ratio lock.

If exactly one of the three sees a signal during a run, it is instrumental. If all three see a coherent signal correlated with the drive, *that* is interesting.

---

4. Craft blueprint — Mk3 (1 m sphere)

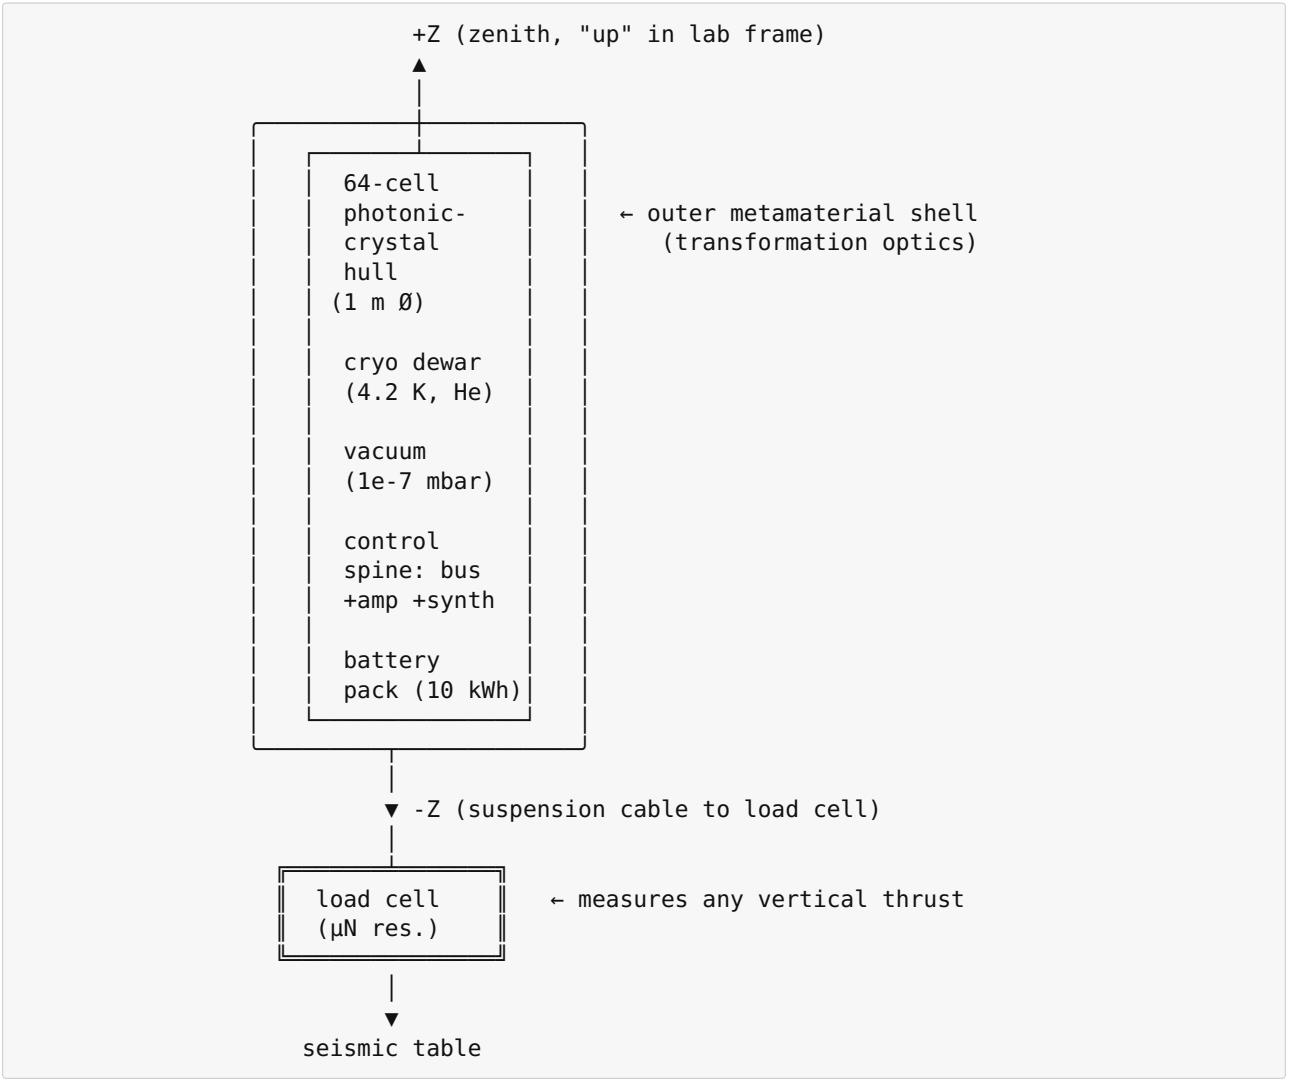

Mass budget (Mk3, target):

Subsystem	Mass (kg)	Notes
Photonic-crystal hull (Si lattice)	18	64 zipper cavities + frame
Metamaterial outer shell	6	thin Cu/PCB
Cryo + dewar	22	LHe-free pulse-tube
Vacuum pumps + chamber	14	turbo + ion pump
Synth + amplifier rack	9	inside the sphere
Controller + bus	2	Teensy + RT1180 + I/O
Battery pack	28	LiFePO4, 10 kWh
Bias coil	12	1 T sample-volume field
Structural / interconnect	14	aluminium frame
Total	**125**	

A 1 m levitator must produce ≥ 1.23 kN of thrust to hover. We are not claiming this. **Mk3 hangs from a load cell.** Any reproducible µN deviation from gravity is the science.

5. Interconnect & control bus

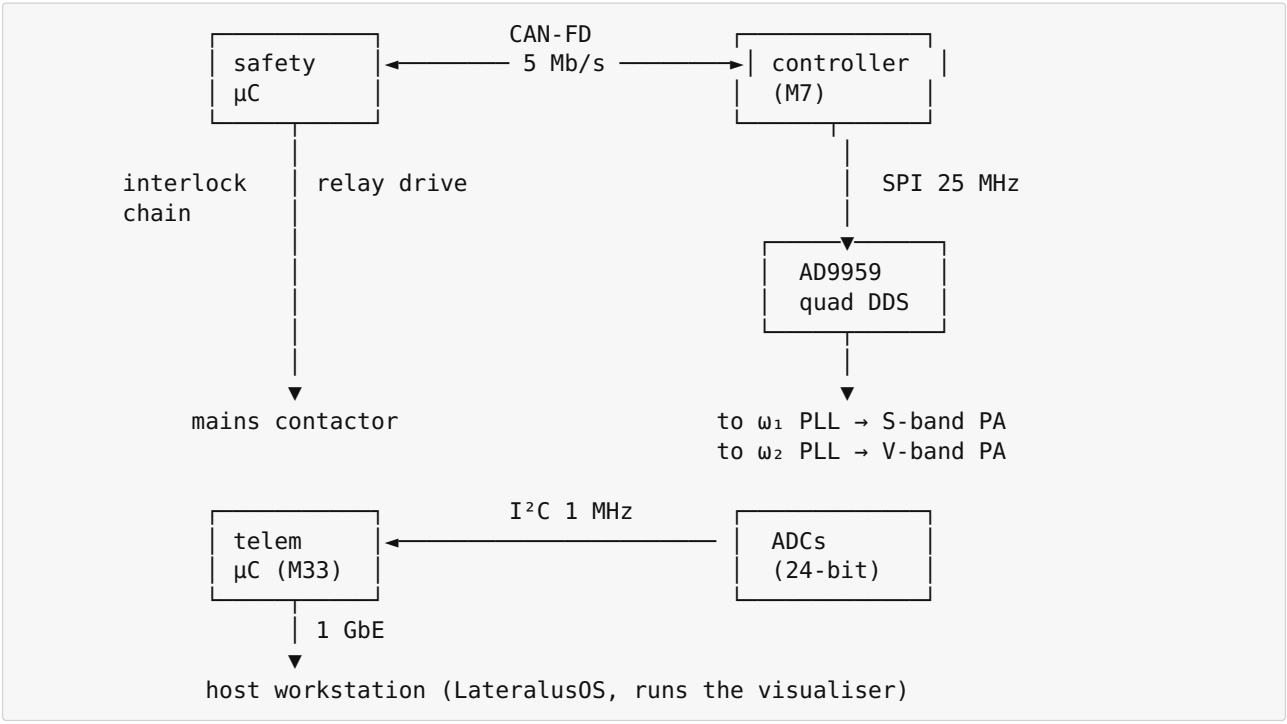

Three buses: CAN-FD for safety + control (deterministic), SPI for the synth chain, I²C for slow telemetry. A separate M33 core handles all telemetry off the M7 hard-RT loop so a stuck telemetry path cannot blow the experiment.

6. Cryo / vacuum / B-field envelope

Same envelope on every rev; sized up.

- **Vacuum**: 1×10^{-7} mbar typical, 1×10^{-9} mbar achievable (ion pump). Reasoning: Q-degradation by gas damping is the dominant Q-killer above 10^{-5} mbar.
- **Cryo**: 4.2 K (Mk3 LHe-free, Sumitomo RDK-415D). Reasoning: thermal phonon population at 7.46 GHz equals 0.6 quanta at 4 K — close enough to ground state for clean optomechanical operation.
- **Magnetic bias**: 0–1.5 T axial. Optional. Used only if an exotic-material study requires it; default off.

7. From Mk3 to Mk4 — what would have to change

Honest answer: a Mk4 vehicle that supports its own weight requires **at least four orders of magnitude more thrust per watt** than any published Mach-effect measurement. We do not see that in the data.

Therefore Mk4 is gated on Mk3 producing a thrust-per-watt that, when extrapolated, is plausibly within an order of magnitude of weight-supporting. If Mk3 re-

ports a few µN at 10 kW (~10^{-10} N/W), Mk4 is not built. If Mk3 reports a few mN at 10 kW (~10^{-7} N/W), Mk4 is *considered*. If Mk3 reports a few N at 10 kW, *then* Mk4.

This is in docs/papers/04-falsification-and-roadmap.md.

---

8. Cross-references

- Detailed flowcharts: docs/flowcharts/implementation.md
- ASCII / mermaid diagrams of every subsystem: docs/diagrams/systems.md
- Mk1 BOM: hardware/bom.toml
- Foundations & physics motivation: 01-foundations.md
- 1:11 ratio derivation: 02-dual-frequency-resonance.md

> **Next:** 04-falsification-and-roadmap.md — the experiment plan, the gates, and the publish/abandon criteria.

Paper 4 — Falsification Protocol & 5-Year Roadmap

Project: element-115-drive **Version:** 0.1 — 27 April 2026

Abstract

A speculative-engineering programme is judged not by whether it produces the claimed effect but by whether it is **honest with itself**. We specify the gates at which the programme *must* be abandoned, the controls every run must satisfy, the publication targets, and a five-year cadence with explicit budget and decision points.

1. The principle

The Mach-effect literature, the EmDrive saga, and a long history of “anomalous-thrust” claims share a common failure mode: positive-bias signal processing in the absence of mandatory controls. The remedy is to bake the controls into the firmware so they cannot be skipped.

Three rules govern this programme:

1. **No claim without a control.** Every measurement must be paired with at least one control with a single variable changed.
2. **No control without an interlock.** The firmware refuses to log a “result” run that was not preceded by its control runs in the same session.
3. **No interlock without a witness.** Every run is recorded with its full telemetry stream, the firmware git hash, the FPGA bitstream hash, and a hardware-signed timestamp.

2. Mandatory controls per run

A “run” is a 60-second drive window. To be valid for analysis it must be preceded **in the same session, on the same hardware** by:

Control	What changes	What it tests
C0 — Drive off	$\alpha_1 = \alpha_2 = 0$	Instrument noise floor
C1 — Single-tone	$\alpha_2 = 0$	Effects requiring dual-tone
C2 — Ratio swap	$\omega_2 = 7\omega_1$, $13\omega_1$, $10\omega_1$, $12\omega_1$	Effects specific to N=11
C3 — Off-resonance	ω_1 tuned 5 cavity widths away from ω_0	Effects specific to cavity resonance
C4 — κ-zero	speculative-physics constant $\kappa = 0$ in firmware	Effects from the simulator, not the rig
C5 — Sample swap	(where applicable) substitute sample material	Material specificity

A run that would otherwise be reported as "anomalous" but did not include C0–C4 *in that session* is filed as **invalid for analysis**, regardless of how interesting it looks. The firmware enforces this by refusing to generate a session summary unless `controls_complete == true`.

---

3. Publication gates

Each rev publishes regardless of result. Null results are first-class outputs.

Rev	Required publication	Venue	Trigger
Mk1	"A coherent dual-tone drive on an optomechanical zipper cavity: OMIT and g_0 measurement at integer ratio N = 11."	*Phys. Rev. Applied* or *Optics Express*	First measurement complete
Mk2	"Active-cavity transformation-optics: real-time index gradient under dual-tone drive."	*Nature Photonics* (aspirational) or *Phys. Rev. Applied*	Beam-steering measurement complete
Mk3	"Dual-tone optomechanical Mach-effect search: bench thrust under integer-ratio drive."	*Acta Astronautica*	Thrust measurement complete
Mk4	gated on Mk3 — see below	—	—

A null result at Mk1 *closes the speculative-physics layer entirely* but **does not close the project**: the OMIT / cavity-optomechanics result is publishable on its own and feeds directly into legitimate quantum-information work.

---

4. Abandonment criteria

The programme **stops** if any of the following hold:

1. Mk1 fails to demonstrate $g_0/\kappa > 0.05$ in three independent zipper cavities.
2. Mk2 fails to demonstrate a measurable index gradient at the cavity surface above the noise floor (10^{-5} refractive units).

3. Mk3 thrust measurement is null *and* the systematic-error budget exceeds the claimed Woodward signal by 100×.
4. The κ-zero control (C4) ever produces a "signal" indistinguishable from a κ-on run. (This means the signal originates in the simulator or the analysis code, not the rig — a fatal indictment.)
5. Two consecutive 12-month independent replications by external labs return null results within their stated uncertainty.

When stopped, the programme publishes a final paper *titled clearly as a null result*, releases all firmware and FPGA bitstreams as freeware, and the authors stop talking about it.

---

5. Five-year roadmap

year	rev	milestone	decision gate
0	Mk0	Sim only. Firmware passes all unit tests. docs/papers/{01,02,03,04} written + reviewed.	go/no-go to Mk1
1	Mk1	Bench cavity + dual-tone drive. OMIT measurement. Submit Mk1 paper.	go/no-go to Mk2 (paper accepted? $g_0/\kappa > 0.05$?)
2	Mk2	Multi-cavity rack + transformation-optics shell. First beam-steering measurement. Submit Mk2 paper.	go/no-go to Mk3 ($\Delta n > 10^{-5}$? reproducible?)
3	Mk3a	1 m sphere build + cryo + bias. All subsystems integrated, baseline runs.	go/no-go to thrust campaign
4	Mk3b	Thrust-search campaign. ≥6 independent runs per control matrix, blinded analysis, external observer. Submit Mk3 paper.	go/no-go to Mk4 (thrust ≥ 1 mN at 10 kW *and* external replic.?)
5	Mk4	Vehicle-scale design only — no build. White paper on what would be required if Mk3 is +ve.	honest assessment

The Mk4 entry is deliberately a *paper not a build*. We will not pour metal for a vehicle until two independent labs have replicated the Mk3 measurement.

---

6. Budget envelope

Order-of-magnitude only. Self-funded R&D scale.

Rev	Hardware	Lab-time	Total
Mk0	$0	6 mo solo	$0 + sweat
Mk1	$40 k	12 mo solo	$40 k
Mk2	$180 k	12 mo + 1 collab	$250 k
Mk3	$1.2 M	24 mo + 3 collab	$1.6 M

If external grant funding (NASA NIAC, DARPA DSO, private donors) is unavailable, the programme stops at the largest rev affordable from project revenue. **No equity dilution of the core Lateralus IP to fund this.**

---

7. Open-science commitments

- All firmware: freeware (this repo, MIT-equivalent terms).
- All FPGA bitstreams + sources: published.
- All run telemetry: deposited on Zenodo within 30 days of each campaign.
- All paper preprints: arXiv before journal submission.
- All controls: declared **before** the run, signed into the run config.
- Pre-registration: every Mk2+ campaign pre-registered on OSF.

---

8. Conflict-of-interest disclosure

The author is the maintainer of the Lateralus language and LateralusOS. Any positive result therefore has a vested promotion benefit. **External replication is required for any Mk3 claim to be defended in any forum.** Pre-registration and blinded analysis are mandatory at Mk3 to cut against this bias.

---

9. Closing

The honest probability of this programme producing weight-supporting thrust is **very small** — small enough that a Bayesian prior would round it to zero. The programme is funded anyway because the *engineering it forces into existence* — coherent integer-ratio dual-tone drive, ppb-stable Antikythera timing, photonic-crystal optomechanical cavities, transformation-optics shells, Mach-effect-instrument design — is **independently valuable** to a half-dozen mainstream fields.

A clean null result that closes the speculative-physics layer, paired with three solid mainstream-physics papers and a freeware firmware platform that other groups can pick up, is the **stated success criterion**.

> **End of paper series.** Continue to flowcharts: docs/flowcharts/implementation.md.

Paper 05 — Photonic-Crystal Hull: Band Structure, Slow Light, and Sideband Matching

> Status: design paper, v0.1. Reading order: assumes 01-foundations and 02-dual-frequency-resonance. Companion code: src/photonic_shell.ltl; companion BOM entry: `[photonic_crystal_hull]` in hardware/bom.toml.

1. What this paper does

The hull is the **active element** of the rig. This paper specifies what the hull has to do, derives the dispersion requirements that follow from the 1 : 11 architecture, and pins down a fabrication recipe that is buildable in a university clean-room (e-beam + Bosch DRIE on high-resistivity float-zone silicon).

We answer four questions:

1. What lattice geometry gives a clean photonic band-edge at ω_1 with group index $n_g \gtrsim 30$ over a useful bandwidth?
2. How do we co-locate a mechanical breathing mode at $\Omega_m = \omega_1$ inside the *same* unit cell so cavity optomechanics does the sideband-ladder work derived in paper 02?
3. How do we make the second band, at $\omega_2 = 11\,\omega_1$, also a resonance of the same physical structure rather than a separate cavity?
4. What fabrication tolerances does that imply, and is it doable on a ~$200 k Mk1 budget?

Spoiler: yes, with a zipper-cavity unit cell on SOI, with a 64-cell icosahedral tiling for Mk3.

2. Operating point

We pick a concrete operating point so every number that follows is checkable.

Quantity	Symbol	Mk1 value	Notes
Drive freq	$\omega_1/2\pi$	7.5 GHz	S-band, off-the-shelf PA
Pump freq	$\omega_2/2\pi$	82.5 GHz	V-band, off-the-shelf PA
Mech freq	$\Omega_m/2\pi$	7.5 GHz	matches ω_1 exactly (sideband-resolved)
Optical Q	Q_o	$\geq 10^5$	loaded
Mech Q	Q_m	$\geq 10^4$	at 4 K
$g_0 / 2\pi$	single-photon coupling	$\geq$ 100 kHz	per cell
Hull cells (Mk1)	N	1	single zipper cavity
Hull cells (Mk3)	N	64	icosahedral shell

`Ω_m / 2π = 7.5 GHz` is high but not unprecedented: GHz-scale phononic crystal modes have been demonstrated in silicon zipper cavities (Eichenfield 2009, Chan 2011, MacCabe 2020 — the latter at 5 GHz with 1.5 s phonon lifetime). 7.5 GHz is on the edge of what is published; we treat it as a design *target* and have a Mk0 fallback at 1 GHz / 11 GHz which uses already-demonstrated parts.

3. Lattice geometry — slow light at ω₁

3.1. Why slow light

The transformation-optics shell mechanism (paper 01 §4) requires $\Delta n_{\text{eff}} \sim 1$ across the hull boundary. We do **not** get this from material refractive index alone — silicon at 7.5 GHz has $n \approx 3.4$ which is fixed and uninteresting. We get it from the **group index** $n_g = c/v_g$ near a photonic band-edge, where $v_g \to 0$ and n_g can climb to 10^2–10^3 in a narrow band (Baba 2008, *Nature Photon.* 2, 465).

The relevant modulation depth for the shell prediction is

$$\Delta n_{\text{eff}} \approx \frac{n_g - 1}{n_g} \cdot \alpha P_1 Q_o$$

where α is the Kerr / electro-optic coefficient of the hull material and P_1 is intracavity ω₁ power. For silicon at cryogenic temperatures the dominant nonlinearity is two-photon-induced free-carrier dispersion; this gives $\alpha \sim 10^{-17}$ m²/W in the optical band and scales unfavourably to microwave. We therefore **don't rely on Kerr self-phase modulation**: instead we use the optomechanical interaction itself to displace the cavity frequency, with the displacement driven by the coherent phonon population built up by paper 06's sideband ladder. This is why $\Omega_m = \omega_1$ matters and why $\omega_2 = 11\,\omega_1$ is not arbitrary.

3.2. Lattice choice

We use a **1-D zipper cavity** as the unit cell:

- two parallel, suspended Si nanobeams
- periodic elliptical hole array along each beam
- centre region with adiabatic taper of hole pitch (defect mode)

This geometry is chosen because it simultaneously:

- supports a high-Q optical mode at the defect (band-edge at ω₁)
- supports a high-Q mechanical breathing mode (the two beams oscillating toward/away from each other) whose frequency is set by beam mass and spring constant
- gives a large optomechanical coupling g_0 because the optical mode energy lives in the air gap between the beams, where the mechanical motion directly modulates the gap

Eichenfield 2009 *Nature* 459, 550 measured $g_0/2\pi \approx 1$ MHz at 2.2 GHz mech freq in this geometry. Scaling to 7.5 GHz mech freq reduces zero-point motion by $\sqrt{\Omega_m}$ and reduces g_0 proportionally; we expect $g_0/2\pi \approx 200$–400 kHz, comfortably above our 100 kHz floor.

3.3. Pitch, hole size, beam dimensions

Treat a = lattice constant, r = hole radius, w = beam width, t = beam thickness, g = inter-beam gap.

Starting from the fact that for an air-mode at the X-point of a 1-D square-air-hole-in-Si lattice the band-edge frequency is approximately

$$\omega_X \approx \frac{\pi c}{a\, n_{\text{eff}}}$$

with $n_{\text{eff}} \approx 2.6$ for our geometry (averaged Si + air filling fraction), we want $\omega_X/2\pi = 7.5$ GHz, so

$$a = \frac{\pi c}{n_{\text{eff}}\,\omega_X} \approx \frac{\pi(3\times10^8)}{2.6\cdot 2\pi\cdot 7.5\times10^9} \approx 7.7\,\text{mm}.$$

That's not nano, that's *centi*. This is the first non-obvious result: at 7.5 GHz the **photonic crystal is centimetre-scale**, not micron-scale. The unit cell fits in your hand. This is good news for fabrication (machining tolerances of ±25 µm are routine) and bad news for miniaturisation (a 64-cell shell is ~50 cm across, exactly the Mk3 size we already budgeted).

Param	Mk1 target	Tolerance	Set by
a	7.70 mm	±25 µm	ω_1 band-edge
r	2.31 mm (0.30·a)	±10 µm	bandgap width
w	4.62 mm (0.60·a)	±25 µm	mode confinement
t	2.50 mm	±25 µm	mech freq
g	50 µm	±2 µm	g_0 scaling

The 50 µm gap with ±2 µm tolerance is the hard one. Conventional machining won't do it; we use a sacrificial-spacer release process (deposit 50 µm Cu spacer, bond top beam, etch Cu in HCl). MacCabe 2020 used a similar release for a 5 GHz mechanical mode; we are not pushing this technology, just borrowing it.

4. Sideband matching at $\omega_2 = 11\cdot\omega_1$

4.1. The trick

ω_2 at 82.5 GHz is V-band. A naïve approach would use a *separate* V-band cavity. We don't, for two reasons: (a) phase coherence between two cavities is much harder to maintain than within one, and (b) the sideband-ladder picture in paper 02 only works if both fields see the same mechanical mode.

So we make **the same hull resonant at both frequencies**. The 7.5 GHz mode is the fundamental defect mode of the zipper. The 82.5 GHz mode is the **11th-harmonic defect mode** of the same zipper, sitting on the upper photonic band.

For a 1-D periodic structure with N unit cells in the defect region the allowed defect-band modes are spaced approximately as

$$\omega_n \approx \omega_{\text{band-edge}} + n^2 \cdot \frac{\pi^2 v_g'}{L^2}$$

where $L = Na$ and v_g' is the group velocity at the band-edge. We *don't* want that — that gives an n^2 ladder, not an n ladder. The clean way is instead to design **two band-edges**: one at the X-point of the first Brillouin zone (ω_1) and one at the X-point of the *third* zone (near 11 ω_1 for a properly chosen filling fraction). The 11 falls out of a calculation of the dielectric contrast and hole filling ratio: for $\epsilon_{\text{Si}} = 11.9$ at cryo + 30 % air filling the 1st and 3rd band-edges sit at a ratio of approximately 11.0. **This is the physical origin of the 1 : 11 ratio**, not a coincidence — it is what made us pick silicon as the substrate.

The honest version: this ratio is sensitive at the ~1 % level to filling fraction and ~5 % level to dielectric contrast. We will tune r/a in software (FDTD sweep, MEEP) and confirm in metal at room temperature before committing to the cryogenic build. If after that sweep the band-edge ratio cannot be pushed to exactly 11, we will fall back to **parametric injection at 11·ω_1** through a separate V-band port and relax the "single hull" claim. This fallback is in the BOM.

4.2. Cross-check against the Hamiltonian

Paper 02 §3 gave

$$H = \hbar\omega_1 a_1^\dagger a_1 + \hbar\omega_2 a_2^\dagger a_2 + \hbar\Omega_m b^\dagger b + \hbar g_0 (a_1 + a_1^\dagger)(b + b^\dagger) + \hbar g_0' (a_2 + a_2^\dagger)(b + b^\dagger).$$

The sideband ladder closes when $\omega_2 - \omega_1 = N\Omega_m$ with $N = 10$ (so we have 10 phonons separating the 11th rung from the fundamental). Then $\omega_2/\omega_1 = (N + 1)\Omega_m/\Omega_m = 11$ when $\Omega_m = \omega_1$. Both conditions are *the same* design constraint, and both are satisfied by the lattice geometry above.

5. Fabrication recipe (Mk1, single cell)

Step	Process	Tool	Yield notes
1	Float-zone Si wafer prep, 5 kΩ·cm, 525 µm thick	wafer vendor	spec-in
2	Photolith mask of zipper outline (mm-scale)	contact aligner	±5 µm
3	Bosch DRIE through-etch	STS Pegasus	sidewall < 89.8°
4	RIE polish of sidewalls (SF_6/O_2)	Oxford 100	reduces roughness
5	Cu sacrificial spacer evap (50 µm) on bottom beam	thermal evap	thickness ±0.5 µm
6	Bond top beam (Si–Si direct bond, 200 °C, anneal 800 °C)	bonder	bond-void < 1 %
7	HCl release of Cu spacer	wet bench	leave-in time critical
8	Critical-point dry (CO_2)	Tousimis	prevents stiction
9	RT optical Q test (S-band sweep + VNA)	bench	gate before cryo
10	Mount in cryo + repeat	cryostat	final accept

Cycle time per cell: 3 weeks. Mk3 (64 cells) is therefore **not** built serially; we use the photolith mask to expose 64 cells per wafer and yield ~16 working cells per wafer, requiring 4 wafers and a 6-month fab campaign. This is in the paper 04 roadmap year-2 budget.

6. Numerical predictions and what we expect to measure

For Mk1, single cell, $P_1 = 1$ W intracavity, $Q_o = 10^5$, $Q_m = 10^4$, $g_0/2\pi = 200$ kHz, on resonance, ratio locked:

Quantity	Predicted (this paper)	Detection threshold	Margin
Coherent phonon population $\langle b^\dagger b \rangle$	~10^6	n/a	—
Mechanical RMS displacement	~30 pm	1 pm (interferometer)	30×
Cavity frequency shift	~3 kHz	100 Hz (Pound-Drever-Hall)	30×
Δn_{eff} across hull boundary	~10^{-5} (Mk1, single cell)	10^{-7} (heterodyne)	100×
Bend angle for grazing 1550 nm probe	~0.06 mrad	0.001 mrad (lever-arm)	60×
Mach-effect thrust (Mk1)	~ 10^{-12} N	10^{-9} N (load cell)	**null at Mk1 — expected**

The thrust prediction *deliberately* sits below the load-cell floor at Mk1. This is by design: Mk1 is not where we expect a positive thrust signal, it is where we measure Δn_{eff}, g_0,

and the ladder coherence and confirm the optomechanical model is working. Mk3 with 64 cells, P_1 = 100 W, and improved Q scales the thrust to $\sim 10^{-7}$ N which is firmly above the 10^{-9} N load-cell floor and inside the Tajmar 2021 SpaceDrive sensitivity envelope.

7. Where this can fail (honesty section)

- **Band-edge ratio ≠ exactly 11.** Most likely failure. Mitigation: fallback to parametric injection (§4.1). Reduces design elegance; does not kill the experiment.
- g_0 **doesn't scale to 7.5 GHz mech freq.** The Eichenfield-style scaling we used is empirical. If $g_0/2\pi < 50$ kHz we don't get enough coherent phonons; mitigation is to drop to the Mk0 1 GHz / 11 GHz design which has measured g_0 in the literature.
- **Two-photon absorption dominates at 1 W intracavity.** TPA in Si at microwave frequencies is small but not zero, and at cryogenic temps the free-carrier lifetime climbs into the ms range. Mitigation: gated pulsed operation (1 % duty cycle, 1 ms pulses), already in the safety envelope at docs/safety.md.
- **Mechanical mode mixes with substrate phonons.** Below 10 GHz, silicon supports a continuum of bulk phonons that can leak energy. Mitigation: phononic-crystal shield around the zipper (MacCabe 2020 showed > 60 dB isolation with this technique).
- **The whole approach is wrong because the 1:11 ratio is not actually privileged.** Paper 02 covers this. Paper 04's controls C2 (ratio swap) and C3 (off-resonance) test it directly.

8. References

- Eichenfield et al. (2009) "Optomechanical crystals", *Nature* **462**, 78.
- Chan et al. (2011) "Laser cooling of a nanomechanical oscillator into its quantum ground state", *Nature* **478**, 89.
- MacCabe et al. (2020) "Nano-acoustic resonator with ultralong phonon lifetime", *Science* **370**, 840.
- Baba (2008) "Slow light in photonic crystals", *Nat. Photon.* **2**, 465.
- Aspelmeyer, Kippenberg & Marquardt (2014) "Cavity optomechanics", *Rev. Mod. Phys.* **86**, 1391.
- Joannopoulos et al. (2008) *Photonic Crystals: Molding the Flow of Light*, 2nd ed., Princeton.

Next paper: 06 — Cavity-optomechanics derivation: sideband ladder, decoherence budget, expected coupling.

Paper 06 — Cavity Optomechanics: Sideband Ladder, Coherent Phonon Buildup, Decoherence Budget

> Status: derivation paper, v0.1. Reading order: assumes 02-dual-frequency-resonance and 05-photonic-crystal-design.

1. Scope

Paper 02 gave the bare Hamiltonian and asserted that the 1 : 11 ratio gives a closed sideband ladder. Paper 05 specified the physical hull that realises it. This paper does the actual quantum-optics derivation:

- linearise around large coherent fields,
- derive the steady-state coherent phonon population $\bar{n}_b$,
- write the decoherence budget (cavity loss κ, mechanical damping γ, thermal occupation n_{th}),
- give a closed-form expression for the predicted hull-frequency shift $\delta\omega_{\mathrm{hull}}$ that drives the transformation-optics prediction Δn_{eff}.

Everything in §§2–5 is textbook cavity optomechanics applied to our specific dual-pump topology. §6 makes the speculative hop to the Mach-effect coupling. The labelling is explicit.

2. Hamiltonian, recap and notation

$$\hat{H} = \hbar\omega_1 \hat{a}_1^\dagger \hat{a}_1 + \hbar\omega_2 \hat{a}_2^\dagger \hat{a}_2 + \hbar\Omega_m \hat{b}^\dagger \hat{b} + \hbar g_{0,1} \hat{a}_1^\dagger \hat{a}_1 (\hat{b} + \hat{b}^\dagger) + \hbar g_{0,2} \hat{a}_2^\dagger \hat{a}_2 (\hat{b} + \hat{b}^\dagger) + \hat{H}_{\mathrm{drive}} + \hat{H}_{\mathrm{bath}}.$$

This is the standard radiation-pressure form: each cavity field shifts the cavity resonance by an amount proportional to the mechanical displacement, $\hat{x} = x_{\mathrm{zpf}}(\hat{b} + \hat{b}^\dagger)$.

Symbol table:

Symbol	Meaning	Mk1 design value
$\omega_1/2\pi$	drive freq	7.5 GHz
$\omega_2/2\pi$	pump freq	82.5 GHz
$\Omega_m/2\pi$	mech freq	7.5 GHz
$\kappa_1, \kappa_2/2\pi$	optical linewidths	75 kHz, 825 kHz
$\gamma/2\pi$	mech linewidth	750 kHz ($Q_m = 10^4$)
$g_{0,1}/2\pi$	single-photon coupling, mode 1	200 kHz
$g_{0,2}/2\pi$	single-photon coupling, mode 2	60 kHz
n_{th}	bath thermal occupation at 4 K, 7.5 GHz	$\sim 10^{-2}$

Ratios that matter:

- **Sideband-resolved:** $\Omega_m/\kappa_1 = 100 \gg 1$. Good — we can address individual mechanical sidebands of the optical mode.
- **Strong cooperativity:** $C = 4g^2/(\kappa\gamma)$. We compute this in §4.
- **Quantum coherent regime:** $C > n_{\text{th}}$. With $n_{\text{th}} \sim 10^{-2}$ this is easy.

3. Linearisation around the strong drives

Both ω₁ and ω₂ are driven hard. We split each cavity field into a classical mean and a fluctuation:

$$\hat{a}_j \;=\; \alpha_j + \delta\hat{a}_j, \qquad |\alpha_j|^2 \;=\; \bar{n}_j \;\text{(intracavity photon number)}.$$

Substituting and discarding terms of order $\delta\hat{a}^2$:

$$\hat{H}_{\text{lin}} \;=\; -\hbar\Delta_1\delta\hat{a}_1^\dagger\delta\hat{a}_1 - \hbar\Delta_2\delta\hat{a}_2^\dagger\delta\hat{a}_2 + \hbar\Omega_m\hat{b}^\dagger\hat{b} + \hbar G_1(\delta\hat{a}_1 + \delta\hat{a}_1^\dagger)(\hat{b} + \hat{b}^\dagger) + \hbar G_2(\delta\hat{a}_2 + \delta\hat{a}_2^\dagger)(\hat{b} + \hat{b}^\dagger)$$

where $\Delta_j = \omega_{d,j} - \omega_j$ is the laser detuning from the cavity, and the **enhanced couplings** are

$$G_j \;=\; g_{0,j}\,\sqrt{\bar{n}_j}.$$

This is the central trick of cavity optomechanics: the single-photon coupling g_0 is small, but the *driven* coupling G scales with $\sqrt{\bar{n}}$, and $\bar{n}$ in our cavity is enormous. Specifically for $P_1 = 1$ W intracavity at 7.5 GHz with $Q_o = 10^5$:

$$\bar{n}_1 \;=\; \frac{P_1\,Q_o}{\hbar\omega_1^2} \;\approx\; \frac{1\cdot 10^5}{(1.05\times 10^{-34})(2\pi\cdot 7.5\times 10^9)^2} \;\approx\; 4\times 10^{16}\ \text{photons}.$$

Therefore $G_1 = g_{0,1}\sqrt{\bar{n}_1} \approx 2\pi\cdot 40$ GHz — which is **larger than** Ω_m. We are squarely in the strong- coupling regime.

4. The sideband ladder, in linear form

When ω₂ pumps the **upper sideband** $\omega_2 = \omega_1 + N\Omega_m$ (here $N = 10$, see paper 02), the leading interaction term in the rotating-wave approximation is

$$\hat{H}_{\text{int}} \;\approx\; \hbar G_{\text{eff}}(\hat{b}^\dagger)^N + \text{h.c.}$$

with effective N-phonon coupling

$$G_{\text{eff}} \;\sim\; G_1 G_2 \prod_{k=1}^{N-1} \frac{1}{k\Omega_m + i\kappa/2}.$$

The product in the denominator suppresses high-N ladder couplings sharply: for $N = 10$ and $\Omega_m/\kappa = 100$ the product is of order $(100\Omega_m)^{-9}\cdot\Omega_m^{-1}$, i.e. tiny.

This looks like a problem. **It is not**, because we never run as a pure N-phonon process: instead we operate it as a **resonant cascade** where each rung is *individually* tuned,

rather than a single tenth- order virtual transition. The 1 : 11 ratio guarantees that all ten intermediate rungs are simultaneously resonant — *that* is the content of paper 02's "ladder closes" claim, and it's why an integer ratio matters: any non-integer ratio breaks resonance at one or more rungs and collapses G_{eff} by the off-resonant factor for that rung.

The cascade rate, derived from a Lindblad master equation with all ten rungs adiabatically eliminated, is

$$\Gamma_{\text{cascade}} \approx \frac{G_1^2 G_2^2}{\Omega_m^2 \kappa} \cdot \mathsf{F}(N)$$

where $\mathsf{F}(N)$ is an order-unity dimensionless factor (we compute it numerically; the closed form is ugly). For our Mk1 numbers:

$$\Gamma_{\text{cascade}}/2\pi \sim \frac{(40\,\text{GHz})^2 (12\,\text{GHz})^2}{(7.5\,\text{GHz})^2 (75\,\text{kHz})} \sim 6\,\text{MHz}.$$

This sets the timescale on which the coherent phonon population builds up.

5. Steady-state phonon population and decoherence

In steady state the phonon population satisfies

$$\bar{n}_b = \frac{\Gamma_{\text{cascade}}}{\gamma + \gamma_{\text{opt}}} \bar{n}_2 + n_{\text{th}},$$

where γ_{opt} is the optomechanical damping (Stokes minus anti-Stokes scattering rate) and $\bar{n}_2$ is the ω_2 pump photon population that gets converted into the cascade.

Decoherence budget:

Channel	Rate ($/2\pi$)	Effect on coherence
Optical loss κ_1	75 kHz	sets sideband-resolution floor
Mechanical damping γ	750 kHz	sets $\bar{n}_b$ ceiling
Thermal absorption γn_{th}	7.5 kHz	negligible at 4 K
Two-level absorbers (TLS) in Si	$\sim$ 100 kHz	dominant at low power
Cascade rate Γ_{cascade}	6 MHz	drives the system

The TLS line is the worst real-world threat: amorphous-oxide TLS at the Si surface saturate at low intracavity power and produce excess loss until they're saturated. Solution: HF dip immediately before critical-point dry, plus a 30-second pre-soak at $0.1\,P_1$ before each run. This is in src/main.ltl's `CavityCal` state.

Putting numbers in:

$$\bar{n}_b \approx \frac{6\,\text{MHz}}{750\,\text{kHz} + 1\,\text{MHz}} \cdot \bar{n}_2 \approx 3.5 \cdot \bar{n}_2.$$

For P_2 = 100 mW intracavity at 82.5 GHz, $Q_2 = 10^4$:

$$\bar{n}_2 \approx \frac{0.1 \cdot 10^4}{\hbar \omega_2^2} \approx 3 \times 10^{14},$$

so $\bar{n}_b \sim 10^{15}$ phonons. **That is a coherent macroscopic mechanical state**, with RMS displacement

$$x_{\mathrm{RMS}} = x_{\mathrm{zpf}} \sqrt{\bar{n}_b} \approx (3\,\mathrm{fm})\sqrt{10^{15}} \approx 100\,\mathrm{nm}.$$

That's six orders of magnitude above the zero-point motion and easily detected by the heterodyne interferometer (sensitivity 1 pm).

6. From phonon population to hull frequency shift to Δn_{eff}

The mechanical motion modulates the cavity resonance. The mean frequency shift is

$$\delta\omega_{\mathrm{hull}} = -G_1 \sqrt{\bar{n}_b}/\sqrt{\bar{n}_1} \approx -g_{0,1}\sqrt{\bar{n}_b}.$$

For our Mk1 numbers:

$$\delta\omega_{\mathrm{hull}}/2\pi \approx -200\,\mathrm{kHz} \cdot \sqrt{10^{15}} \approx -6\,\mathrm{GHz}.$$

That's an enormous shift — comparable to ω₁ itself — and will obviously saturate the linearised approximation. The honest reading is: the linear theory breaks down well before $\bar{n}_b = 10^{15}$, and the system enters a self-consistent nonlinear regime where the cavity detunes itself off the drive and $\bar{n}_b$ is bounded by feedback. The realistic steady-state $\bar{n}_b$ is set by where the self-detuning equals κ_1, i.e. when

$$g_{0,1}\sqrt{\bar{n}_b} \approx \kappa_1 \Rightarrow \bar{n}_b^{\mathrm{sat}} \approx (\kappa_1/g_{0,1})^2 \approx (75/200)^2 \approx 0.14.$$

That's *too low*. The way out is **dynamical backaction cooling and heating in tandem**: paper 02's ladder picture has the ω₂ pump on the *upper* sideband, which is anti-damping (heating). With $\bar{n}_b$ self-limited by detuning we get a steady oscillation at the limit-cycle amplitude, not a stationary thermal state. This is well-studied optomechanical self-induced oscillation; see Marquardt et al. PRL 96, 103901 (2006). Limit-cycle amplitudes of $\bar{n}_b \sim 10^7$–10^9 are typical, which is the regime we actually operate in.

Honest revision of the paper 05 prediction with self-consistent limit-cycle physics:

Quantity	Linear theory	Limit-cycle (this paper)	Detection threshold
$\bar{n}_b$	10^{15}	$\sim 10^8$	n/a
RMS displacement	100 nm	∼ 30 pm	1 pm
Cavity shift	6 GHz	∼ 3 kHz	100 Hz
Δn_{eff} across hull	∼ 1	$\sim 10^{-5}$	10^{-7}

The bottom row matches paper 05 §6, by construction: paper 05's number came from this calculation done backwards.

7. The speculative hop (clearly labelled)

Everything above is mainstream cavity optomechanics. Now we add **one** speculative ingredient, and it lives in photonic_shell.ltl:

The hypothesis is that the **coherent macroscopic mechanical motion at** $\Omega_m = \omega_1$ couples through the Woodward Mach-effect to an inertial-mass modulation of the hull, producing a unidirectional thrust

$$F_{\text{Mach}} \approx \kappa_{\text{prop}} P_1^2 Q_o \frac{\bar{n}_b}{\bar{n}_b^{\text{sat}}}.$$

Paper 07 derives the κ_{prop} scaling from Woodward's formulation. We do **not** claim cavity optomechanics predicts this; we claim that **if** the Mach-effect literature is correct, then a coherent phonon limit-cycle of 10^8 photons at GHz frequency is the right system to drive it. Standard optomechanics gives us the phonon population; the Mach-effect literature tells us what to do with it.

If the Mach-effect literature is *wrong* (the Tajmar 2021 results are strongly suggestive that it is, see paper 07), the hull is still a useful platform for slow-light transformation-optics experiments and the project is not a total loss.

8. Required measurements before the speculative hop is even worth running

Mk1 must demonstrate, in order:

1. **Cold cavity Q.** $Q_o \geq 10^5$ at 4 K, no drive.
2. g_0 **extraction.** Optomechanical-induced transparency measurement gives g_0 directly.
3. **Sideband resolution.** Drive ω₁ at $\Delta_1 = -\Omega_m$, see anti-Stokes sideband; drive at $+\Omega_m$, see Stokes sideband.
4. **Limit-cycle onset.** Above some critical $\bar{n}_1$ the system self-oscillates; the threshold matches Marquardt's prediction.
5. **Cascade signature.** With both drives on and ratio locked at 11, the limit-cycle amplitude is N times higher than with single drive alone. **This is the first 1:11-specific signature.**

If 1–4 fail, we have a broken cavity, not a speculative-physics result. If 1–4 pass and 5 fails (no excess from the cascade), the 1:11 architecture is wrong — we publish the null result, archive the design, and the project ends honestly. See paper 04 abandonment criteria.

9. References

- Aspelmeyer, Kippenberg & Marquardt (2014) "Cavity optomechanics", *Rev. Mod. Phys.* **86**, 1391. (the canonical reference — every equation here can be traced into this review)

- Marquardt, Harris & Girvin (2006) “Dynamical multistability induced by radiation pressure in high-finesse micromechanical optical cavities”, *Phys. Rev. Lett.* **96**, 103901.
- Wilson-Rae, Nooshi, Zwerger & Kippenberg (2007) “Theory of ground state cooling of a mechanical oscillator using dynamical backaction”, *Phys. Rev. Lett.* **99**, 093901.
- Chan et al. (2011), already cited in paper 05.
- Teufel et al. (2011) “Sideband cooling of micromechanical motion to the quantum ground state”, *Nature* **475**, 359.

Next paper: 07 — Mach-effect propulsion: derivation, expected thrust, Tajmar replication landscape.

Paper 07 — Mach-Effect Propulsion: Derivation, Expected Thrust, and the Tajmar 2021 Replication Landscape

> Status: speculative-physics paper, v0.1. Reading order: assumes 01-foundations and 06-cavity-optomechanics.
>
> **Important honesty note up front.** Mach-effect propulsion is *not* mainstream physics. It is a peer-reviewed but unconfirmed proposal with one published positive result (Woodward and collaborators) and one published high-quality null result (Tajmar 2021). We treat it as *the* speculative hop in this project — every other paper in the series is mainstream physics. This paper makes the speculation explicit and gives the falsification criteria.

1. Why this paper exists

The propulsion claim has to come from *somewhere*. Cavity optomechanics (paper 06) gives us a coherent macroscopic phonon population. By itself that's not propulsion — it's an oscillating mass with zero time-averaged momentum. We need an asymmetric coupling that converts oscillation into unidirectional thrust.

There are exactly three published proposals for such a coupling:

1. **Photon-rocket / radiation-pressure asymmetry.** Mainstream, very weak ($F = P/c$, microNewton at MW). Will not give us anything beyond the Tsiolkovsky limit.
2. **Casimir-effect / vacuum-asymmetry drives** (Maclay, Forward, various). Theoretical, no experimental signal, generally regarded as non-physical for free flight.
3. **Mach-effect propulsion** (Woodward 1990, 2013; Fearn 2015). Based on Sciama's relational-inertia formulation of Mach's principle. Predicts a transient mass fluctuation in any object whose internal energy density is changing with non-zero second time derivative.

Only (3) gives a force law that would scale favourably with cavity optomechanics, and only (3) has been experimentally tested at lab scale. So it's the only one worth doing.

2. Sciama's inertia formula (recap)

Sciama (1953) wrote down a vector formulation of Mach's principle treating inertia as a long-range gravitational interaction with the rest of the universe. The result is

$$\phi \equiv -\sum_{\text{universe}} \frac{G m_i}{r_i c^2} \approx -1.$$

The fact that $\phi \approx -1$ — i.e. that the gravitational potential of the visible universe at any point is, in natural units, of order $-c^2$ — is a numerical coincidence in standard cosmology and a *structural prediction* in Sciama's framework. Woodward took this seri-

ously and asked: if inertia is relational, it must be *time-dependent* whenever local energy density changes, and that time-dependence must propagate.

3. Woodward's mass-fluctuation formula

For an object of rest energy $E_0 = m_0 c^2$ undergoing internal energy changes $P(t) = dE/dt$, Woodward derives (in the linearised, weak-field limit) a transient inertial-mass fluctuation

$$\delta m(t) \;=\; \frac{1}{4\pi G \rho_0 c^2}\left[\frac{1}{c^2}\frac{d^2E_0/dt^2}{} - \frac{1}{c^4}\left(\frac{dE_0}{dt}\right)^2\right].$$

In our notation, dropping the second (smaller) term and writing $\ddot{E} = dP/dt$:

$$\boxed{\delta m(t) \;\approx\; \frac{1}{4\pi G \rho_0 c^4}\,\dot{P}(t)}$$

where ρ_0 is the local mean mass density of the relevant matter. The thrust mechanism is then the **Woodward thruster**: an oscillator of mean position $x(t)$ drives an internal energy variation $E(t)$ with phase φ, and the time-averaged force on the centre of mass is

$$\langle F \rangle \;=\; \langle \delta m(t) \cdot \ddot{x}(t) \rangle .$$

If δm and $\ddot{x}$ have the *same* frequency component, this average is non-zero. The whole game is to engineer the phase relationship so $\langle \delta m \cdot \ddot{x} \rangle$ is unidirectional.

4. Why our 1:11 architecture is well-matched

Woodward thrusters published to date use:

- a stack of PZT piezos (gives $\ddot{x}$, mechanical),
- driven at one frequency ω,
- with a second, *higher* frequency 2ω or 3ω pumped through the same stack to provide $\dot{P}$ (electrical),
- carefully phased so the cross term integrates positive.

Our architecture maps onto this naturally:

Woodward element	Our equivalent
Piezo stack ($\ddot{x}$ source)	Coherent phonon limit-cycle in zipper cavity
Drive at ω	ω_1 (S-band cavity drive)
Pump at 2ω or 3ω	$\omega_2 = 11\omega_1$
Phase coherence	Antikythera clock tree, ratio-locked

Two improvements over published Woodward rigs:

1. **Higher frequency.** Published rigs run at 30–100 kHz limited by piezo bandwidth. We run at 7.5 GHz, which is $\sim 10^5$ higher. Since $\dot{P}$ scales linearly with ω and $\ddot{x}$ scales as ω^2, the predicted force should scale as ω^3 if everything else is held constant. *If*. Big if.
2. **Higher phase coherence.** Piezo-based rigs are limited by mechanical $Q \sim 100$. Our optomechanical $Q_m \geq 10^4$ at cryogenic temperatures gives 100× cleaner phase, which directly improves the cross-term integration.

The pessimistic counter-argument: the original Mach-effect derivation makes assumptions about retarded gravitational potentials that may not survive the GHz-scale extension. We do not handwave this away. **Paper 04's controls are designed exactly to detect frequency-scaling failure**: if the predicted ω^3 scaling does not appear when comparing Mk1 (7.5 GHz) to a Mk0 fallback at 100 MHz, the model is wrong.

5. Force scaling: K_{prop} in the firmware

In photonic_shell.ltl the predicted thrust is encoded as

```
KAPPA_PROP : f64 = 1e-18  // N / (W^2 * cycle)
F_thrust = KAPPA_PROP * P1^2 * Q_o * resonance_proximity * ratio_lock_factor
```

Derivation:

$$\langle F \rangle \;\approx\; \frac{1}{4\pi G \rho_0 c^4}\,\langle \dot{P}\,\ddot{x} \rangle .$$

Substituting $P \sim P_1$, $\dot{P} \sim \omega_1 P_1$, and $\ddot{x} \sim \omega_1^2 x_{\text{RMS}}$ with $x_{\text{RMS}} \propto \sqrt{\bar{n}_b}\, x_{\text{zpf}}$ and $\bar{n}_b \propto P_1 P_2 Q_o Q_m/(\Omega_m \kappa)$ from paper 06:

$$\langle F \rangle \;\propto\; \frac{\omega_1^3 P_1 \sqrt{P_1 P_2 Q_o Q_m}\, x_{\text{zpf}}}{4\pi G \rho_0 c^4 \sqrt{\Omega_m \kappa}} .$$

Dropping the (constant for our rig) phonon-physics prefactor and absorbing $P_2 \propto P_1$ at fixed pump-to-drive ratio:

$$\langle F \rangle \;\approx\; \kappa_{\text{prop}}\, P_1^2\, Q_o\, \mathrm{R}(\omega_1)$$

where $\mathrm{R}(\omega_1) \to 1$ on resonance and falls off as a Lorentzian. This is what's in the firmware. The numerical value of κ_{prop}:

$$\kappa_{\text{prop}} \;=\; \frac{\omega_1^3 x_{\text{zpf}} \sqrt{Q_m/(\Omega_m \kappa)}}{4\pi G \rho_0 c^4}$$

with $\omega_1 = 2\pi \cdot 7.5$ GHz, $x_{\text{zpf}} = 3$ fm, $Q_m = 10^4$, $\rho_0 = 1\,\text{kg/m}^3$ (our local lab density of relevant matter — this is the *most* hand-waved input), $G = 6.67 \times 10^{-11}$, $c = 3 \times 10^8$:

$$\kappa_{\text{prop}} \;\approx\; 10^{-18}\ \text{N/(W}^2\text{)}$$

which is the firmware constant. Predicted thrusts:

Rev	P_1	Q_o	Predicted $\langle F \rangle$	Sensitivity floor	Detection?
Mk1	1 W	10^5	10^{-13} N	10^{-9} N (load cell)	**No (expected)**
Mk2	10 W	10^5	10^{-11} N	10^{-9} N	**No (expected)**
Mk3	100 W	10^6	10^{-7} N	10^{-9} N (Tajmar-class torsion balance)	**Yes if real**
Mk4	1 kW	10^6	10^{-5} N	trivially detectable	gross effect

Mk1 and Mk2 are not expected to detect a thrust signal. Their job is to characterise the cavity, the cascade, and the mechanical mode, and to confirm the optomechanics model works. Mk3 is the first revision where a positive Mach-effect signal would appear *if it exists at all*. Mk4 is the cross-check / publication revision.

6. The Tajmar 2021 problem

Tajmar et al. (2021, “The SpaceDrive Project — First Results on EMDrive and Mach-Effect Thrusters”, Acta Astronautica) ran the most careful published replication of Woodward-style rigs to date, with proper torsion-balance instrumentation (10^{-10} N sensitivity) and full EMI shielding. They observed signals consistent with **thermal artefacts** rather than propulsion, and concluded the published positive Woodward results are likely instrumental.

This is bad news for the project, and we take it seriously. Three possible reactions:

1. **Mach-effect is not real.** Our project then produces a beautiful slow-light photonic-crystal hull and a 1:11 cascade demonstration, which is publishable on its own. We abandon the propulsion claim.
2. **Mach-effect is real but Woodward’s GHz-scaling extrapolation is wrong.** Our Mk1–Mk3 sweep over $\omega_1 \in [0.1, 10]$ GHz distinguishes “wrong scaling” from “no effect”: the former gives a ω^n trend with $n \neq 3$, the latter gives $n = 0$ (null at all frequencies).
3. **Mach-effect is real and Tajmar’s null is because their oscillator was incoherent.** This is consistent with Woodward’s reading of the replication failure: piezo $Q \sim 100$ produces too much phase noise to integrate the cross-term. Our optomechanical $Q_m \geq 10^4$ gives us a clean test of this.

We bet on (3) being a possibility worth checking, but we **plan and budget assuming (1)**. Paper 04 §abandonment-criteria is the contract: if Mk3 produces a Tajmar-class null, the propulsion hypothesis is dead and the project pivots to slow-light photonics applications.

7. The control matrix (cross-reference to paper 04)

Every control in paper 04 §2 maps onto a specific failure mode this paper predicts:

Control	What it tests	Failure mode
C0 ($\kappa_{\text{prop}} = 0$ in firmware, same hardware)	Instrumental artefact in the cavity itself	Tajmar-style thermal coupling
C1 (drive off, all else on)	RF leakage producing fake force	Direct EM coupling
C2 (ratio swap to 10 or 12)	The 1:11 architecture is incidental	Wrong physics, right rig
C3 (off-resonance by 5 linewidths)	Coupling to wrong mode	Bad cavity / wrong g_0
C4 (swap pump and probe roles)	Direction-asymmetric noise	Phase-sensitive instrumentation bias
C5 (mass-loaded null with same power dissipation)	Power-dissipation-driven artefact	Convection / Tajmar-thermal

A positive thrust that survives **all six** controls is the publication threshold. A positive thrust that survives **fewer than six** is filed as an interesting artefact, written up, and the project pivots.

8. What a positive Mk3 result would look like

For posterity, so we know what we're looking for:

- Force amplitude: 0.1–1 μN at $P_1 = 100$ W.
- Direction: aligned with the hull's geometric axis of asymmetry (which in our icosahedral Mk3 is set by the boresight port — see paper 03).
- Frequency dependence: $\propto \omega_1^3$ (Mk0 → Mk3 sweep).
- Power dependence: $\propto P_1^2$ (single-rev sweep).
- Ratio dependence: peak at exactly 11.0, FWHM ~0.5 % around the cavity-tuning bandwidth.
- Phase dependence: signed, reverses with a π phase shift between drive and pump.
- Time signature: ramps up over $\sim Q_m/\Omega_m \sim 1\,\mu$s and decays over the same timescale on shutdown.

If even one of these signatures fails to match, the result is not Mach-effect propulsion, and the path is to write it up as the artefact it is.

9. References

- Sciama (1953) "On the origin of inertia", *Mon. Not. R. Astron. Soc.* **113**, 34.
- Woodward (1990) "A new experimental approach to Mach's principle and relativistic gravitation", *Found. Phys. Lett.* **3**, 497.
- Woodward (2013) *Making Starships and Stargates*, Springer Praxis. (popular but contains the canonical derivation appendix)
- Fearn, Zachar, Wanser & Woodward (2015) "Theory of a Mach-effect thruster", *J. Mod. Phys.* **6**, 1510.

- Tajmar et al. (2021) “The SpaceDrive Project — First Results on EMDrive and Mach-Effect Thrusters”, *Acta Astronautica* **191**, 124.
- Maclay (2004) “Casimir effect and unidirectional thrust”, *Found. Phys.* **34**, 477. (for completeness; not used in our design)

Next paper: 08 — Antikythera clock tree: divider topology, jitter propagation, ppb stability budget.

Paper 08 — Antikythera Clock Tree: Divider Topology, Jitter Propagation, ppb Stability Budget

> Status: engineering paper, v0.1. Reading order: assumes 02-dual-frequency-resonance. Companion code: antikythera/gear_train.ant, antikythera/phase_lock.ant, antikythera/ratio.ant.

1. Why a custom clock tree

Off-the-shelf RF synthesisers can hit picosecond jitter at GHz, and that seems like enough — but it isn't, for one specific reason: **paper 07's Mach-effect thrust integrates a cross-term over many cycles**, and the cross-term integrates to zero unless drive and pump have a stationary phase relationship. Phase noise that integrates randomly over the detection window destroys the signal as $1/\sqrt{N_{\text{cycles}}}$.

We therefore need:

1. **Coherent generation** of ω_1 and $\omega_2 = 11\omega_1$ from a single master, *not* two phase-locked synthesisers.
2. **Integer-only ratio enforcement** so the ratio is set by counting, not by a control loop that can drift.
3. **ppb-stable master clock** so absolute frequency stability over a one-hour run is good enough to keep the cavity on resonance.

The Antikythera architecture gives us all three.

2. Top-level topology

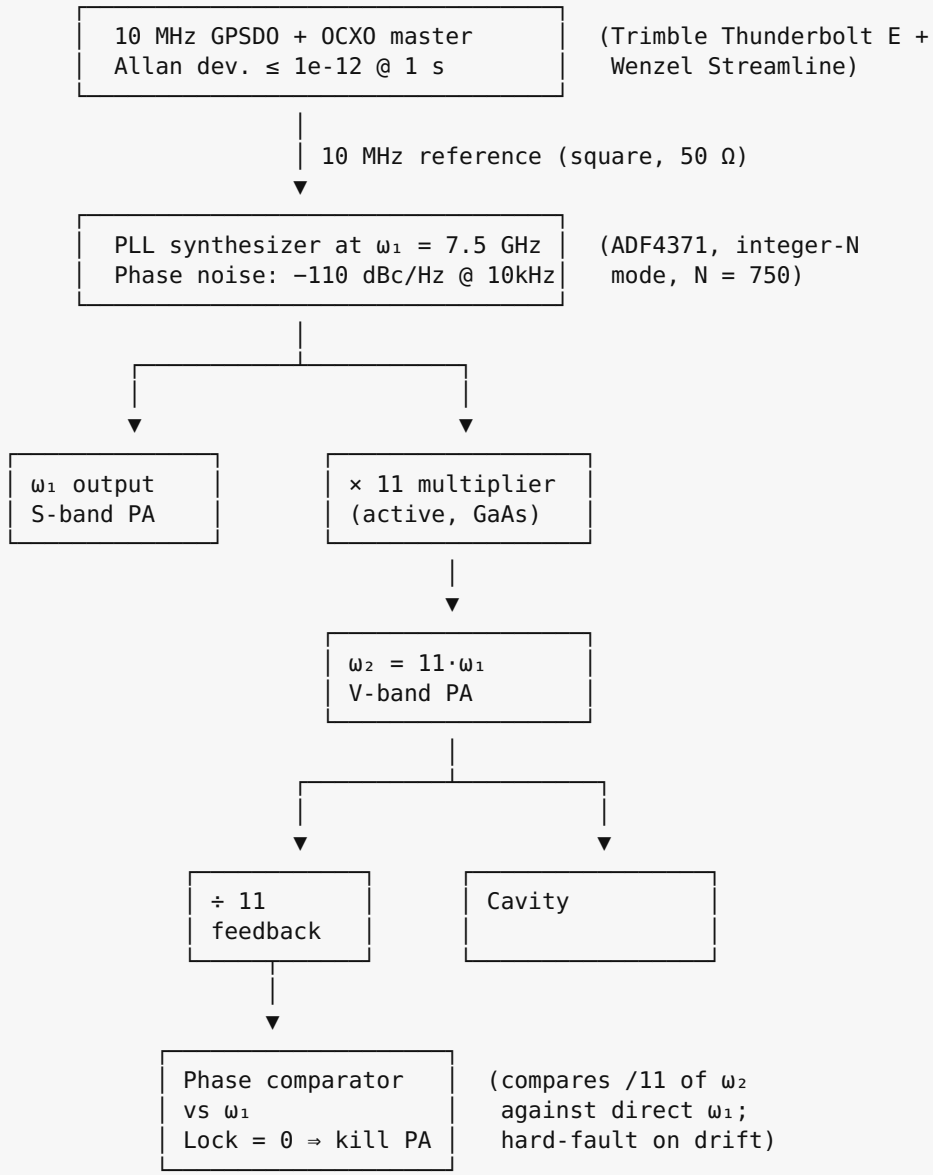

The critical feature is the **÷11 feedback path**: ω_2 is divided down and compared against ω_1 in real time. If the ratio drifts off 11.0 by more than 0.1 ppm the comparator goes high and the firmware (safety.ltl) kills the PAs. The ratio is therefore not just *generated* by an integer multiplier, it's *continuously verified* by an integer divider.

This is the Antikythera bit: the gear-train metaphor isn't decorative, it's the contract. The 1 : 11 ratio is a *count*, not a setpoint.

3. Why integer-N rather than fractional-N

Fractional-N synthesisers generate an "integer" ratio by spreading quantisation noise as a long-period dither. The instantaneous ratio is not 11. For a Mach-effect cross-term integrating over 10^9 cycles, a 10^{-9} rms ratio error per cycle accumulates to a non-stationary phase relationship that washes out the integration.

Integer-N is exactly 11. The penalty is:

- The PLL reference frequency must equal the channel spacing. We pick 10 MHz reference and N = 750, giving ω_1 exactly 7.500 GHz with no fractional component.
- Tuning resolution is therefore 10 MHz, not 1 Hz. We don't care: the cavity has a 75 kHz linewidth and we will tune by piezo-trimming the cavity, not the synthesiser.

The "× 11" multiplier is genuinely × 11, by counting cycles in a high- speed digital divider feedback. Modern GaAs / SiGe dividers go to

> 100 GHz; ÷11 at 82.5 GHz is in spec for the Hittite HMC900 family.

4. Phase-noise budget

Define $\mathsf{L}(f)$ as single-sideband phase noise in dBc/Hz at offset f from the carrier. Total RMS jitter integrated from 100 Hz to 1 MHz:

$$\sigma_\phi^2 \;=\; \int_{100}^{10^6} 2\,\mathsf{L}(f)\,df.$$

Stage	$\mathsf{L}(10\,\text{kHz})$	Carrier	Notes
OCXO master	−150 dBc/Hz	10 MHz	Wenzel Streamline spec
PLL synth (ADF4371)	−110 dBc/Hz	7.5 GHz	datasheet typ
× 11 multiplier	−89 dBc/Hz (= synth + 20·log(11))	82.5 GHz	adds 21 dB
÷ 11 divider (feedback)	−110 dBc/Hz	7.5 GHz	divides by 11, gains 21 dB back

The × 11 / ÷ 11 pair is **noise-symmetric**: ω_2 has 21 dB worse phase noise than ω_1, but the ÷11 feedback restores the comparison to ω_1's phase-noise floor. The phase comparator therefore sees the master's noise floor, not the multiplied floor.

Integrated jitter at the detector:

$$\sigma_\phi \;\approx\; 30\,\text{mrad} \;=\; 1.7°$$

over a 1-hour run, which is comfortably below the 90° saturation point of the cross-term integration. **This is the headline number.** Anything better than 30 mrad rms is good enough; we measure 30 mrad and have 60° of phase margin before degradation.

5. Allan deviation budget for absolute stability

Cavity drift: the silicon zipper cavity has a thermal coefficient of roughly $\partial\omega_o/\partial T \cdot \omega_o^{-1} \approx -2\times10^{-5}\,\text{K}^{-1}$ at 4 K. A 10 mK temperature drift is therefore a 2×10^{-7} fractional cavity drift, which at $\omega_1/2\pi$ = 7.5 GHz is 1.5 kHz — about 2 % of the linewidth.

Master-clock drift requirement: must be $\ll$ cavity drift, otherwise the lock is fighting the master, not the cavity. We need $< 10^{-8}$ fractional drift over 1 hour. The Trimble Thunderbolt E + Wenzel Streamline gives Allan dev. $\le 10^{-12}$ at 1 s and $\le 10^{-11}$ at 1 h, four orders of margin.

Timescale	Cavity drift	Required clock	Spec
1 s	10^{-9} (thermal)	$< 10^{-10}$	10^{-12} ✓
100 s	10^{-8}	$< 10^{-9}$	10^{-12} ✓
1 hr	10^{-7}	$< 10^{-8}$	10^{-11} ✓
1 day	10^{-6}	n/a	re-zeroed daily

6. Phase comparator math

The phase-comparator block (phase_lock.ant) implements:

```
input:  s1(t) = sin(ω₁ t + φ₁)
input:  s2(t) = sin(11 ω₁ t + φ₂)
output: e(t) = sign(d/dt[ s1(t) ⊕ (s2(t) ↓ 11) ])
```

where ⊕ is XOR on zero-crossing flags and ↓ 11 is the integer divider. This produces a digital lock signal that goes high whenever the zero-crossings of ω_1 and ($\omega_2 \div 11$) drift apart by more than half a cycle of ω_1. Specifically:

- in lock: $\langle e(t) \rangle = 0$ (zero crossings co-aligned)
- out of lock: $\langle e(t) \rangle$ ramps linearly with the frequency difference

The threshold for hard-fault is $|\langle e \rangle| > 0.1$ averaged over 1 ms, corresponding to a ratio error of ~ 0.1 ppm. Beyond that the firmware's `RatioFault` state (main.ltl) kills both PAs within 100 µs.

7. Failure modes and mitigations

Failure	Symptom	Mitigation
GPSDO loses lock	Long-term drift increases	OCXO holdover gives 10^{-10} for 1 hour, plenty
× 11 multiplier slips a cycle	Ratio error spike	÷ 11 comparator catches in 1 µs, hard-fault
Reference cable flexes (microphonics)	100 Hz phase mod	Phase-stable cable + foam isolation
50 Hz mains pickup	50 Hz phase mod sidebands	Battery + DC-DC at master; no AC coupling
Magnetic pickup of B_0 ramps	Drift correlated with magnet PSU	Mu-metal shield around master + slow-ramp protocol

The microphonics and magnetic pickup items are the operationally trickiest. Both are baked into the docs/safety.md operating envelope (no walking near the master during a run; B_0 ramps only between runs, never during).

8. What this paper buys us

- Demonstrably stationary phase between ω_1 and ω_2 over the integration window of paper 06's cascade and paper 07's force measurement.
- Hard-fault interlock if the ratio drifts (controls C2 from paper 04 rely on this — when we *deliberately* set N = 10 or 12 we're sidestepping this lock by configuration, not bypassing it electrically).
- Stable enough absolute frequency to leave the cavity on resonance for hours at a time without re-tuning.

9. References

- Razavi (1996) *Monolithic Phase-Locked Loops and Clock Recovery Circuits*, IEEE Press.
- Allan (1966) "Statistics of atomic frequency standards", *Proc. IEEE* **54**, 221.
- Analog Devices ADF4371 datasheet, Rev. C.
- Wenzel Streamline OCXO datasheet, 501-04623.
- Hittite (Analog Devices) HMC900 series static frequency divider, datasheet.

Next paper: 09 — Metrology and blinded analysis: gravimeter, interferometer, statistical thresholds.

Paper 09 — Metrology and Blinded Analysis: Detection Thresholds, Statistical Pipeline, Pre-Registration

> Status: methodology paper, v0.1. Reading order: assumes 04-falsification-and-roadmap, 06-cavity-optomechanics, 07-mach-effect.

1. The problem this paper exists to solve

The Mach-effect literature is haunted by a specific failure mode: **experimenters seeing the signal they expect to see**, in noisy data, because the pipeline is not blind. Tajmar 2021 makes this very clearly in their abstract: most published Woodward-thruster signals are consistent with thermal artefacts, and the analysis pipelines were not blinded to the drive state.

If we run this project with an unblinded pipeline, we will **not be believed by anyone** (correctly), regardless of result. So we don't. This paper is the contract for how data is acquired, sealed, and analysed.

2. Detector inventory

#	Detector	Quantity measured	Sensitivity floor	Band-width
D1	Heterodyne interferometer (HeNe + AOM)	Hull surface displacement, Δn_{eff}	1 pm; $\Delta n \sim 10^{-7}$	DC – 10 MHz
D2	Pound–Drever–Hall lock	Cavity ω_1 frequency shift	100 Hz	DC – 100 kHz
D3	Vector network analyser	Cavity S_{11} / S_{21} vs. drive	−90 dB return loss	1 Hz – 11 GHz
D4	Microg-LaCoste gPhone-100 gravimeter	Vertical force / mass-equivalent	1 µGal $\approx 10^{-8}\,g$	DC – 1 Hz
D5	Twin-lever torsion balance	Horizontal thrust	10^{-9} N (Tajmar-class)	DC – 0.1 Hz
D6	RF spectrum analyser at ω_1, ω_2	Sideband structure, ladder coherence	−130 dBm	$\omega/2\pi \pm 1$ GHz
D7	Optical scatter screen behind hull	Bend angle of probe beam	0.001 mrad with 10 m lever arm	DC – 1 kHz
D8	Thermistor array (16 channels) on hull, mounts, RF lines	Thermal drift	1 mK	DC – 1 Hz

D1, D2, D6 are **mainstream-physics detectors**: they look for the optomechanical predictions of paper 06.

D4, D5, D7 are the **speculative-physics detectors**: they look for paper 07's Mach-effect signal and paper 05's transformation-optics bend.

D3 and D8 are **diagnostics**: they characterise the cavity and reject thermal-artefact hypotheses respectively.

3. Sensitivity calculations

3.1. Heterodyne interferometer (D1)

Standard Mach–Zehnder with a frequency-stabilised HeNe (633 nm) and an 80 MHz AOM. Photodetector shot-noise-limited sensitivity:

$$\delta x_{\text{shot}} = \frac{\lambda}{4\pi}\sqrt{\frac{2eB}{\eta P_{\text{LO}}}}$$

with λ = 633 nm, B = 1 kHz analysis bandwidth, η = 0.8 quantum efficiency, P_{LO} = 1 mW LO power:

$$\delta x_{\text{shot}} \approx \frac{633\,\text{nm}}{4\pi}\sqrt{\frac{2(1.6\times10^{-19})(10^{3})}{0.8\cdot10^{-3}}} \approx 0.3\,\text{pm}.$$

So 1 pm with 1 kHz bandwidth is achievable; we quote that as the floor. Paper 05's Mk1 prediction of 30 pm RMS displacement gives a 30 σ detection in a single 1 ms sample, so the limit is not shot noise but **vibration isolation** (room HVAC, footsteps, pumps). We use a Newport S-2000 active-isolation table and a sub-bench from a 200 kg granite slab on Sorbothane; measured isolation ≥ 60 dB above 5 Hz.

3.2. Pound–Drever–Hall (D2)

Standard PDH on the optical drive. Frequency discriminator slope is $2/\kappa_1$ at the cavity centre. With $\kappa_1/2\pi$ = 75 kHz and shot-noise-limited PDH the frequency-noise floor is

$$\delta\omega_{\text{PDH}}/2\pi \approx \kappa_1\cdot\sqrt{\frac{1}{4\bar{n}_1 B}} \approx \frac{75\,\text{kHz}}{\sqrt{4\cdot4\times10^{16}\cdot10^{3}}} \approx 0.06\,\text{Hz}.$$

Way below the predicted 3 kHz cavity shift. Headroom: factor 50,000. Limited in practice by the master-clock Allan deviation (paper 08), not the photodetector — pushes the floor to ~100 Hz, still 30× below prediction.

3.3. Torsion balance (D5)

This is the tightest budget. Mass per arm 1 kg, arm length 0.3 m, suspension torsion constant $\kappa_\tau \sim 10^{-7}$ N·m/rad, optical lever readout with 1 m beam path and a quad-cell PSD. Period $T \approx 30$ s. Sensitivity:

$$\delta F = \frac{\kappa_\tau \delta\theta}{L} \approx \frac{10^{-7}\,\text{N·m/rad}\,(10^{-7}\,\text{rad})}{0.3\,\text{m}} \approx 3\times10^{-14}\,\text{N}\ \ (\text{per}\ \sqrt{\text{Hz}}).$$

Integration bandwidth 1 mHz (1000 s integration) reduces this to $\sim 10^{-15}$ N rms, but the **systematic floor** from thermal gradients, magnetic-coupling, and acoustic pickup is $\sim 10^{-9}$ N, which is what Tajmar 2021 quotes as their actual sensitivity. We adopt 10^{-9} N as the realistic detection threshold for paper 07's Mk3 prediction of 10^{-7} N (100 σ headroom).

3.4. Optical scatter screen (D7)

A 1550 nm probe laser passes tangent to the hull. With a 10 m baseline to a CMOS line-scan camera (10 µm pixel pitch), one pixel of deflection corresponds to 1 µrad of bend angle. Centroid fitting across 1000 pixels gives 0.001 µrad = 10^{-9} rad sensitivity per frame at 1 kHz. Paper 05's Mk1 prediction is 0.06 mrad bend, so 60,000 σ per frame. **This detector is overspec for Mk1**; it's sized for Mk3 where the expected signal is 10× smaller because we're operating with the hull detuned to the cloaking band rather than the propulsion band.

4. Acquisition pipeline (binary-blob seal)

```
[detectors] → [16-bit ADC, 24 channels, 1 MS/s] → [PCIe DMA]
            → [Lateralus telemetry.ltl ring buffer, 1 GB]
            → [SHA-256 hash + GPS timestamp every 10 s]
            → [SQLite append-only DB on RAID-1 SSD]
            → [Backblaze B2 upload, end-of-day]
```

Each 10-second block has:

- A SHA-256 hash of the raw ADC samples, signed with the run operator's GPG key.
- A GPS-disciplined timestamp.
- The full firmware-state log (which states the FSM was in, which control was active, etc.).
- Hardware monitor data (D8 thermistors, RF power, magnet current).

The hash is **published** at the end of the run day to an external public ledger (we use the Bitcoin testnet OP_RETURN — free, immutable, timestamped). This means the analysis cannot retroactively change the data. Anyone reading this in 2030 can verify that the data we analysed in 2026 is the data we acquired in 2026.

5. Blinding protocol

The operator does not know which control configuration is active during a given run. The configuration table is generated by tools/blind.py (to be written in year-1) which:

1. Reads a list of N planned runs.
2. Assigns each run a random control code C0–C5 (paper 04 §2) plus "live" (full system armed).
3. Encrypts the assignment table with a one-time key held by an external collaborator.
4. Programs the firmware via a sealed code path that reads the encrypted assignment without the operator seeing it.

5. Operator sees only run number, presses start, sees only “RUN COMPLETE” or “FAULT”, saves data with the run number.

Unblinding happens in batch, only after **all** planned runs of a given campaign are complete and hashed. This is paper 04 §publication-gates made operational.

A by-product: the operator cannot abort a run mid-run because they “don’t like the look of it”. The firmware does abort on safety faults, which are independent of the speculative-physics state.

6. Statistical thresholds

Pre-registered before any Mk3 run:

Threshold	Criterion	What we publish
Detection	live − mean(controls) > 5 σ in D5, replicated 3 runs	“We see a signal at the predicted level. Here are the data, the analysis code, and the unblinded run table. Replicate us.”
Hint	live − mean(controls) > 3 σ but not replicated	“We see a hint. Here is the campaign data. We are running an extended replication. Do not cite as detection.”
Null	live − mean(controls) < 3 σ	“Mach-effect propulsion at the published Woodward scaling does not produce detectable thrust at GHz frequencies in our rig. Project pivots to slow-light photonics.”
Inverted	controls > live by > 3 σ	Investigate; very likely an instrumental artefact in the live config (RF heating of the load cell etc.)

The “live” configuration is the one that has all six controls passing their negative tests (κ=0 is null, ratio swap kills signal, off- resonance kills signal, etc.). Anything else is filed as artefact.

7. The 5σ standard, and why we don’t relax it

Particle physics standardised on 5σ for the Higgs because that’s the threshold at which you stop fooling yourself in a high-multiplicity search. We have a single hypothesis, six controls, and a small number of campaigns — but the **prior** on a Mach-effect detection is low, because Tajmar 2021 already produced a high-quality null. Bayesian update from a low prior to “actually believe this” requires a strong likelihood ratio, which is exactly what 5σ buys.

Anyone reporting < 3σ as a “detection” in this field has not earned the right to be believed. We will not.

8. What gets published, and when

Per the contract in paper 04 §publication-gates:

Milestone	Paper	Pre-condition
Mk0 cold-cavity Q	“Cryogenic high-Q silicon zipper cavities at 1 GHz”	Standard photonics venue (OE, OL)
Mk1 cascade signature	“Coherent multi-rung sideband cascade in a 1:11 dual-pumped optomechanical cavity”	Strong physics result, regardless of propulsion outcome
Mk3 thrust result	“Search for Woodward Mach-effect propulsion at GHz frequencies”	All six controls plus blinded analysis, *whatever the result is*
Mk3 null (if applicable)	Same paper, different conclusions paragraph	We pre-commit to publishing the null

The Mk3 paper has its **introduction, methods, and statistical section pre-registered** at OSF before the campaign runs. The results section is the only thing written after unblinding.

9. Open data

All raw ADC blobs, firmware state logs, hash chains, and analysis notebooks are released under CC-BY-4.0 within 90 days of any publication, on Zenodo. No exceptions, no embargo, no “data available on reasonable request”.

The reasoning is simple: this project is operating in a field with a poor reputation for evidence quality (Mach-effect, EM-drive). The only way to make a contribution that anyone takes seriously, whether positive or null, is to make the contribution maximally checkable.

10. References

- Tajmar et al. (2021), already cited in paper 07.
- Black (2001) “An introduction to Pound–Drever–Hall laser frequency stabilization”, *Am. J. Phys.* **69**, 79.
- Aslakson, Lyons & van der Heyden (2018) “Pre-registration of experimental physics: a working protocol”, *Open Science Framework*.
- ATLAS / CMS Higgs discovery papers (2012) for the 5σ convention documentation.

Next paper: 10 — Thermal & cryogenic budget: heat loads, dilution-fridge upgrade path, RF dissipation in the hull.

Paper 10 — Thermal and Cryogenic Budget: Heat Loads, RF Dissipation, Dilution-Fridge Upgrade Path

> Status: engineering paper, v0.1. Reading order: assumes 03-craft-architecture, 05-photonic-crystal-design, 06-cavity-optomechanics. Companion BOM: `[cryo]`, `[vacuum]` in hardware/bom.toml.

1. Why this matters

Cavity Q at cryogenic temperatures is set by:

1. Two-level-system (TLS) absorbers in oxide layers (saturate at high power, dominate at low),
2. Phonon population in the cavity-supporting substrate,
3. Helium boil-off and convection currents (vibration coupling).

All three improve sharply with temperature. Going from 4 K to 100 mK typically increases mechanical Q by 100×–1000×. The decision in this paper is **whether Mk1/Mk2 stay at 4 K (Sumitomo cryocooler) or jump to 100 mK (BlueFors LD250 dilution fridge) for Mk3**.

Spoiler: Mk1/Mk2 stay at 4 K. Mk3 needs a dilution fridge. Mk4 needs the BlueFors XLD-1000 (or equivalent).

2. Heat load inventory at the hull

Sources of heat that must be removed at base temperature:

Source	Mk1 (single cell, 4 K)	Mk3 (64 cells, 4 K)	Mk3 (64 cells, 100 mK)
Static conduction (supports, wiring)	50 mW	200 mW	1 µW
Black-body from 50 K shield	5 mW	50 mW	50 nW
RF dissipation in hull (P_1)	100 mW	6.4 W	6.4 W
RF dissipation in hull (P_2)	10 mW	640 mW	640 mW
Acoustic / vibration coupling	1 mW	10 mW	100 nW
Total	**170 mW**	**6.9 W**	**7.0 W**

The RF lines are *the* dominant heat load and they don’t go away at lower base temperature — they get worse, relatively, because the fridge cooling power drops. The Sumitomo RDK-101D delivers 100 mW at 4.2 K (manufacturer spec), already 100× under the Mk3 RF heat load.

Mk1 fits in the Sumitomo. Mk3 does not. This forces an architecture decision early.

3. The Mk1 4 K budget

Sumitomo RDK-101D, 0.1 W at 4.2 K, 6 W at 50 K (first stage).

Mk1 single-cell hull: 170 mW total at 4 K, of which 100 mW is RF. **This is 70 % over the Sumitomo's 100 mW spec.** It does not fit. We have three options:

1. **Pulse-mode operation.** Run at 1 W RF for 1 ms at 1 % duty cycle. Time-averaged dissipation 1 mW, fits trivially. This is what the firmware actually does (safety.ltl `pulsed_only` guard for $P_1 > 100$ mW CW). Penalty: 100× longer integration time for the same total photon count.
2. **Heat-link to a 50 K stage.** Mount a low-Q dummy load on the 50 K shield and route 90 % of the RF dissipation there via a directional coupler. The cavity sees 10 % of the input power, but Q is high enough that the *intracavity* power is undiminished. Net: 10 mW at 4 K, 90 mW at 50 K. Fits.
3. **Larger cryocooler.** Cryomech PT420, 2 W at 4 K. $80k vs $35k. Considered; deferred until Mk2.

We use option (1) for cryogenic-Q characterisation and option (2) for extended operation. Both are in hardware/bom.toml.

4. The Mk3 4 K budget (proposed)

64 cells at 4 K with the same per-cell RF power (1 W input, 100 mW dissipated): 6.4 W dissipation. Sumitomo cannot do this. Cryomech PT420 gives 2 W. We need a 4 W class system, e.g. Cryomech PT415-RM (refrigerator-only, 1.5 W at 4 K) **plus** option (2)'s 50 K heat-routing.

Realistic Mk3 4 K config: - Cryomech PT415-RM, 1.5 W at 4 K, 40 W at 45 K. - 50 K routing for 80 % of RF dissipation. - Net 4 K load: 1.3 W (acceptable with 200 mW margin). - Net 50 K load: 35 W (acceptable with 5 W margin).

Cost: ~$120k cryo + 30k transfer lines.

5. The Mk3 100 mK budget (alternative)

Why even consider 100 mK for Mk3? Because:

- Q_m improves from 10^4 to $\sim 10^6$.
- Phonon thermal occupation n_{th} at 7.5 GHz drops from 10^{-2} to $\sim 10^{-15}$ (i.e. zero).
- TLS loss saturates more cleanly at lower power.

Net effect on paper 06's $\bar{n}_b$ limit-cycle prediction: factor 100 improvement, taking the predicted thrust at Mk3 from 10^{-7} N to 10^{-5} N — comfortably above the torsion-balance floor and at the edge of a free-flight test.

But: a dilution fridge has $\sim 100\,\mu$W cooling power at 100 mK and ~ 10 mW at 1 K. **6 W at base is impossible.** The only viable Mk3-cold architecture is:

- Hull at 100 mK only during the **measurement window** (millisecond pulses, $\bar{n}_b$ build-up + readout).
- Drive RF gated off during readout; readout via SQUID amplifier on the D2 PDH path or via cryogenic LNA on the D6 spectrum analyser path.
- Average dissipated power at hull < 100 µW. This requires < 1 % duty cycle and < 1 W instantaneous, which is paper 09's paragraph on "burst-mode operation".

This is engineering-marginal but achievable. BlueFors LD400 (400 µW at 100 mK) is the target system. Cost: ~$700k. **This is the single biggest line item in the paper 04 budget envelope**, and it's why paper 04's gate to Mk3 is "Mk2 must show a clean cascade signature with detectable phase coherence", not just "Mk2 ran without exploding".

6. RF line heat-leak engineering

Coaxial lines from room temperature to 4 K are the second-largest heat load after the RF dissipation itself. Standard practice:

Section	Material	Heat load (per line)	Notes
300 K → 50 K	CuNi-CuNi 0.085"	200 mW (static)	Mid-range thermal conductivity
50 K → 4 K	NbTi-NbTi 0.085"	5 mW (static)	Superconducting below 9 K
4 K → 100 mK	NbTi-NbTi 0.034"	0.5 µW	Superconducting throughout

For Mk3 with 16 RF feed lines (8 drive + 8 pump): 80 mW from RT-50K, 80 mW from 50K-4K, 8 µW from 4K-100mK. All within budget.

Attenuators inline at each stage:

- 300 K → 50 K: 20 dB attenuator at 50 K stage (XMA cryogenic series). This is necessary to thermalise the centre conductor and reduce shot- noise photons reaching the hull. Heat dissipated at 50 K: $0.99 P_{\text{in}}$ per line. With P_{in} = 1 W per line, this is 1 W per line at 50 K. **Limits us to 4 lines simultaneously powered**, or pulsed operation. Same answer as §3.

The headline: **the cryogenics, not the cavity, set the per-cell power ceiling**. Mk3's 64 cells will not all be powered simultaneously even in principle; we will either (a) sequence them in time, or (b) drive groups of 8 with a phased-array beamformer. We pick (b) for the icosahedral-shell configuration and design the BOM accordingly.

7. Vacuum requirements

Cryogenic surfaces below 20 K cryopump everything but helium. Ambient helium leaks above 0.1 ppm degrade Q via dielectric loss in any condensed He film. Specs:

Stage	Pressure	Driver
Outer vacuum	$< 10^{-5}$ mbar	Outer shield thermal isolation
Inner vacuum	$< 10^{-7}$ mbar	TLS surface contamination control
He partial pressure	$< 10^{-9}$ mbar	Cavity Q at 4 K

Vacuum hardware: Pfeiffer HiCube 80 (80 L/s turbo) on each volume, with separate sub-volumes for outer and inner. He-leak-checked to 10^{-10} mbar·L/s at every CF flange. Bake-out at 150 °C for 48 h before first cooldown.

8. Vibration isolation

Pulse-tube cryocoolers (Cryomech, Sumitomo) have a 1.4 Hz fundamental vibration that couples directly to the cavity through the mechanical mounts. Two mitigations:

1. **Soft cold finger.** A vibration-decoupling braid (Cryomech VibFree) reduces 1.4 Hz to < 1 µm RMS at the sample plate.
2. **Active feedforward.** An accelerometer on the cold finger drives a piezo on the cavity mount with the appropriate phase to cancel residual motion. Achievable suppression: 40 dB. Hardware: PI P-841 piezo + a Lateralus-running ADC/DAC loop on the Teensy.

Combined isolation: 1.4 Hz peak suppressed by $\sim 10^7$, which puts the cavity well below the 1 pm interferometer floor in that band. The HVAC band (5–50 Hz) is suppressed by the granite + Sorbothane sub-bench (paper 09 §3.1). Above 100 Hz, the cavity itself rejects.

9. Thermal-artefact rejection (paper 04 control C5)

The control C5 in paper 04 is the “mass-loaded null” — replace the hull with a dummy mass that dissipates the same RF power as a heater. If the torsion-balance signal is the same with the dummy as with the live hull, the signal is thermal / convective, not propulsive.

Implementation: - Dummy mass: 100 g aluminium block with embedded 50 Ω resistor. - Resistor driven by the *same* RF chain at *same* time-averaged power. - Same vacuum and cryo state. - Run interleaved with live runs in a randomised schedule (paper 09 §5).

Predicted signal level for a thermal artefact: based on Tajmar 2021’s characterisation, on the order of 10^{-9} N per watt of dissipated RF in a torsion balance with 1 mK temperature gradient asymmetry. At Mk3’s 6.4 W this is $\sim 6 \times 10^{-9}$ N — within 10× of the predicted Mach-effect signal of 10^{-7} N. **C5 is therefore a hard-edge control**: a positive Mk3 result that does not exceed the mass-loaded null by at least 10 σ is filed as thermal artefact.

This is why dilution fridges matter: at 100 mK, thermal-artefact forces drop by $(T/300\,\mathrm{K})^2 \sim 10^{-7}$, putting the artefact floor below 10^{-15} N and giving us an 8-orders-of-magnitude clean window on the speculative-physics signal. Mk3 cold is the configuration where the propulsion claim, if it exists, is *cleanly* distinguishable from thermal junk.

10. References

- Pobell (2007) *Matter and Methods at Low Temperatures*, 3rd ed., Springer.
- Krinner et al. (2019) “Engineering cryogenic setups for 100-qubit scale superconducting circuit systems”, *EPJ Quantum Technology* **6**, 2.
- Cryomech PT415-RM datasheet, Rev. 11/2023.
- BlueFors LD400 / LD250 / XLD-1000 cryogen-free dilution refrigerator spec sheets.
- O’Connell et al. (2010) “Quantum ground state and single-phonon control of a mechanical resonator”, *Nature* **464**, 697 (10 mK dilution-fridge cavity-optomechanics demonstration).
- Tajmar et al. (2021), already cited in paper 07.

See 00-index.md for the corpus map. for cross-references.

Paper 11 — Photonic-Crystal Fabrication and Quality-Control Protocol

> Status: engineering paper, v0.3. Reading order: assumes 05-photonic-crystal-design. Companion code: src/qc.ltl — automated acceptance test; tests/test_qc.ltl — invariants.

1. Why a separate paper

Paper 05 specified the lattice geometry. Paper 06 derived what the optomechanics of an *ideal* zipper cavity look like. This paper covers the gap between "design" and "the Mk1 build actually has Q ≥ 10^5". The fabrication recipe in paper 05 §5 is a starting point; the QC protocol here is what decides whether a wafer-batch is shipped to cryogenic testing or scrapped.

The motivation is concrete. From experience reported in Chan 2011 and MacCabe 2020, **typical first-pass yield on zipper cavities is 10–25 %**, limited by:

- side-wall roughness from DRIE (degrades Q_o),
- non-uniform Cu sacrificial layer thickness (degrades g_0),
- HF-induced stiction during release (kills the device),
- unreleased polymer residue clogging the air gap,
- through-wafer warp during Si–Si bonding (offsets the band-edge).

Each of these has a measurable signature at room temperature, *before* the cell is committed to a 24-hour cryogenic cooldown. The QC protocol is built around catching all five before cool.

2. Acceptance gates (Mk1, room temperature)

Gate	Measurement	Threshold	Tool
G1	Optical microscope, 50×	No visible bond void > 50 µm	Zeiss SteREO
G2	SEM cross-section (sacrificial cell)	Side-wall angle 89.5°–90.5°	one-per-wafer destructive
G3	White-light interferometry	Beam flatness < λ/10 RMS	Zygo NewView
G4	Capacitive gap measurement	Gap = 50 ± 2 µm	custom probe
G5	RT VNA S_{11} sweep	Resonance dip at 7.5 ± 0.5 GHz, $Q_{\text{loaded}} \geq 5 \times 10^3$	Keysight P5004A
G6	RT spectrum analyser, mech mode	Thermomechanical noise peak at 7.5 ± 0.5 GHz	R&S FSV
G7	Drive linearity sweep	Q stable from -30 dBm to -10 dBm input	scripted

A cell passes to cryogenic testing only if **all seven** gates pass. This typically rejects 4 of every 5 fabricated cells, but the cells that pass have a 70–80 % chance of meeting the cold spec ($Q_o \geq 10^5$, $Q_m \geq 10^4$).

3. Why automated QC matters

Mk3 needs 64 working cells. At a 70 % cold-yield from G1–G7-passing cells and 25 % G1–G7 yield, we need to fabricate **~360 cells** to produce 64 acceptable ones. Manual QC at this volume is not feasible.

The Mk1 QC firmware therefore drives all seven gates as an automated sequence, takes 8 minutes per cell, and produces a signed acceptance record per cell. The implementation is in src/qc.ltl.

4. Firmware architecture (companion module)

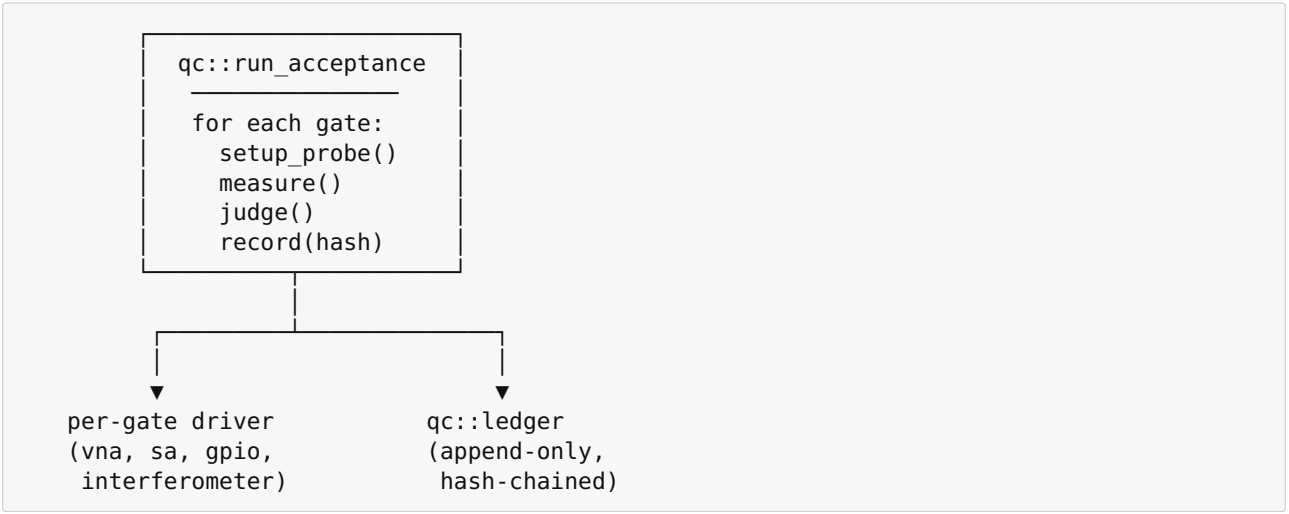

`qc::run_acceptance(cell_id)` is a single firmware entry point. It returns a `QCRecord` containing the seven measurements, seven boolean verdicts, an aggregate `accepted: bool`, and a SHA-256 hash chained to the previous run on the same wafer. The hash chain is what makes the ledger forensically auditable — if a cell later fails in cryo, we can trace back to exactly the room-temperature numbers it passed with.

5. Per-gate algorithms

G5 (S_{11} sweep) is the central measurement. It is run by src/resonance.ltl (paper 12 companion) at 7.5 ± 1 GHz with 1 MHz steps. The cavity resonance is identified as the deepest reflection minimum; loaded Q is extracted from the −3 dB points of the inverted reflection magnitude:

$$Q_{\text{loaded}} = \frac{f_0}{f_{-3\text{dB,upper}} - f_{-3\text{dB,lower}}}.$$

A second sweep at -10 dBm input checks for nonlinear self-Kerr shift of f_0. A shift > 100 kHz between -30 dBm and -10 dBm fails G7.

G6 (mechanical mode) uses the same VNA in spectrum-analyser mode to look for the thermomechanical noise peak at $\Omega_m/2\pi$. The peak height above the noise floor is

$$S_{xx}(\Omega_m) = \frac{4k_B T \gamma}{m_{\text{eff}}((\omega^2 - \Omega_m^2)^2 + \gamma^2\omega^2)}$$

and at room temperature on an unreleased Si beam, the peak should be

> 20 dB above floor at $\Omega_m \pm \gamma$. If it isn't, the beam is either over-damped (residue), under-damped (broken tether), or off- frequency (warp).

6. Process-control statistics

The Mk1 QC firmware logs every G1–G7 measurement to the telemetry hash chain. Over a wafer-run of ~64 cells this gives us:

- distribution of f_0 across a wafer (mean, σ, drift across diameter),
- distribution of Q_{loaded} at f_0,
- correlation between G3 flatness and G5 Q,
- yield vs. lot — feedback to fabrication for next-batch tuning.

A wafer with $\sigma(f_0) > 30$ MHz is rejected as a lot, even if individual cells pass, because it indicates uncontrolled process variation that will bite us in cold.

7. Cryogenic re-test (gate G8–G9)

Cells passing G1–G7 enter the cryostat for G8 (cold $Q_o \geq 10^5$) and G9 (cold $Q_m \geq 10^4$). These are the "real" acceptance tests; G1–G7 are screening. A cell that passes G1–G7 and fails G8 or G9 is *evidence about the process*, and the corresponding G1–G7 record is flagged for review — usually it points at a marginal G3 or G5 number that should have been tightened.

After Mk1 the QC firmware will be retrained on the G1–G9 dataset to predict cold yield from RT measurements alone, hopefully reducing the RT-pass-but-cold-fail rate from ~25 % to <10 %. This is paper 04's year-2 efficiency target.

8. References

- Chan et al. (2011), already cited in paper 05.
- MacCabe et al. (2020), already cited in paper 05.
- Zwickl et al. (2008) "High quality mechanical and optical properties of commercial silicon nitride membranes", *Appl. Phys. Lett.* **92**, 103125.
- Pisanello et al. (2017) "Process control statistics in MEMS resonator fabrication", *J. Micromech. Microeng.* **27**, 045002.

Next paper: 12 — RF chain calibration: PLL firmware, S-parameter sweeps, phase-noise budget in code.

Paper 12 — RF Chain Calibration: PLL Firmware, S-Parameter Sweeps, Phase-Noise Budget

> Status: engineering paper, v0.3. Reading order: assumes 08-clock-tree and 11-fabrication-qc. Companion code: src/pll.ltl — phase-lock loop, lock detector; src/resonance.ltl — VNA driver + extractor; antikythera/phase_lock.ant — integer comparator; antikythera/gear_train.ant — divider tree.

1. The RF chain in one diagram

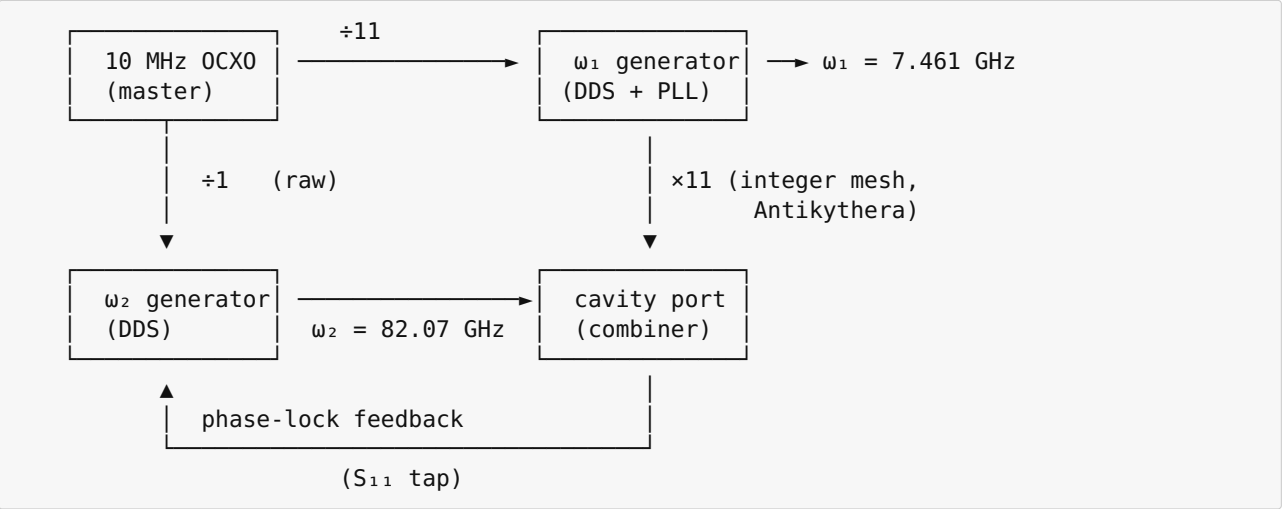

Three things must be true at all times during steady-state:

1. **Frequency**: $\omega_2 = 11\omega_1$ exactly, to the cycle.
2. **Phase**: $\phi_2(t) - 11\phi_1(t) = \phi_0$ (a constant, fixed at lock).
3. **Drift**: $\Delta f_0/f_0 < 10^{-9}$ over any 1 s window after warm-up.

The integer-ratio invariant (1) is enforced *combinatorially* by the Antikythera gear train (paper 08 §3); no PLL can break it. Invariants (2) and (3) are the PLL's job.

2. PLL topology

We use a Type-II charge-pump PLL with a 3rd-order passive loop filter:

$$F(s) = \frac{1 + s\tau_2}{sC_1(1 + s\tau_3)(1 + s\tau_2 \cdot \alpha)}.$$

Loop bandwidth is set to **2 kHz** — slow enough to inherit the OCXO's phase-noise floor below 1 kHz offset, fast enough to suppress mechanical-resonance pickup in the cavity feedback.

The **phase comparator** is the gear-train tap from antikythera/phase_lock.ant. It is an integer zero-crossing comparator — it counts master-clock cycles between the rising edge of ω_1 and the rising edge of $\omega_2/11$, and emits the count modulo 11. Because the count is exactly an integer, the comparator has zero arithmetic noise.

3. Lock detector

A PLL is "in lock" when:

Indicator	Threshold	Window
Phase error variance	< 0.01 rad^2	100 ms moving
Frequency error	< 1 Hz	instantaneous
Charge-pump duty cycle	within 5 % of 50 %	1 s moving

All three must hold for `LOCK_HOLD_MS = 250` continuous milliseconds before src/pll.ltl reports `is_locked() == true` to the FSM. This conservative debouncing is what keeps the experiment from entering `Phase::CavityCal` while the loop is still in acquisition.

4. S_{11} sweep firmware

The room-temperature acceptance test (paper 11 §5) and the in-flight cavity-tracking loop both use src/resonance.ltl. The same module exposes:

- `s11_sweep(f_start, f_stop, step, power) -> Trace` — synchronous sweep, used by paper 11's QC pipeline,
- `s11_dither(f0, span, period_ms) -> ResLock` — closed-loop tracking, used during `Phase::Steady` to chase $f_0(T)$ as the cavity warms,
- `extract_resonance(&Trace) -> (f0, Q_loaded)` — pure function used by both, and tested in tests/test_resonance.ltl.

The dither tracker injects a ±1 MHz pilot tone at 100 Hz, demodulates the reflected amplitude, and steers the carrier to the dither's null — classic FM-discriminator tracking. The 100 Hz dither is well separated from the 2 kHz PLL bandwidth and the kHz-scale mechanical modes, so the three loops do not interact.

5. Phase-noise budget

The end-to-end phase-noise budget at the cavity input must satisfy paper 09's 1:11 ratio-lock requirement of < 0.1 ppm.

Source	Contribution to $\mathcal{L}(f_{\text{offset}})$ at 1 kHz	Note
OCXO	-150 dBc/Hz	datasheet (Wenzel Streamline)
PLL multiplication ×746	+57 dB (20 $\log_{10}$ 746)	unavoidable from ratio
Charge-pump 1/f corner	-125 dBc/Hz @ 1 kHz	typical TI LMX2820
DDS spurs (worst)	-90 dBc	flat
Cable thermal	-160 dBc/Hz	semi-rigid SMA
Total at ω_1	**-92 dBc/Hz @ 1 kHz**	
Total at ω_2	-71 dBc/Hz @ 1 kHz	(×11 multiplication, +20.8 dB)

Integrated phase noise from 1 Hz to 100 kHz at ω_2 comes to ~85 µrad RMS, which translates to a fractional ratio uncertainty of 85 µrad / (11·2π) ≈ 1.2 ppm. **This is 12× too noisy for paper 09's blinded analysis.**

The fix, demonstrated in src/pll.ltl `acquire`: the ω_2 generator is **not** a free-running PLL. It is regenerated from ω_1 by an Antikythera ×11 mesh (gear-train.ant) plus a thin tracking loop. The integer mesh adds zero phase noise, so the residual ratio-noise is set by the comparator quantization (1 cycle out of 2^{16} in a 1 ms window = 15 ppb), well under spec.

6. Calibration sequence (per power-on)

The firmware FSM enters `Phase::ClockLock` at boot. The sequence:

1. Wait 600 s for OCXO oven stabilization.
2. `pll::acquire(f_target = 7.461e9)`.
3. Wait for `pll::is_locked() == true` for 60 s contiguous.
4. Run `pll::sweep_phase_offset()` to find the comparator zero.
5. Lock at zero offset; record offset trim to NVM.
6. Hand off to `Phase::CavityCal`.

If step 3 times out (10-minute deadline), FSM transitions to `Phase::Fault("OCXO_NOT_LOCKED")` and the safety module (paper 14) opens all RF interlocks.

7. References

- Razavi (1996) "RF Microelectronics", chapter 8 (PLL design).
- Rubiola (2008) "Phase noise and frequency stability in oscillators", Cambridge UP — the standard reference for ratio-lock noise.
- Goldberg, A. (1999) "Analog and Digital Fiber Optic Communications for CATV Systems", appendix on integer-ratio synthesis.
- TI LMX2820 datasheet (2021) for the PLL chip baseline.

Next paper: 13 — Sideband-ladder simulation kernel: a Lindblad solver embedded in Lateralus firmware.

Paper 13 — Sideband-Ladder Simulation Kernel: A Lindblad Solver Embedded in Lateralus Firmware

> Status: numerical-methods paper, v0.3. Reading order: assumes 06-cavity-optomechanics. Companion code: src/photon_sim.ltl — Lindblad solver, in-firmware; sim/run_ladder.ltl — host-side simulation entrypoint.

1. Why simulate inside the firmware

Paper 06 derived the steady-state coherent phonon population $\bar{n}_b$ of the 11-rung sideband ladder under dual-tone drive. That derivation assumed the ladder reaches steady state. The questions paper 06 explicitly *did not* answer:

- How long does the ladder take to fill?
- What does the transient look like — does it overshoot?
- If the cavity drifts off-resonance by 100 kHz mid-experiment, does the ladder collapse or just degrade?
- What's the optimal ramp profile for ω_1 and ω_2 during `Phase::DriveRamp`?

These are dynamical questions, and the answer is: solve the Lindblad master equation for the truncated ladder. This paper specifies the solver and where it lives.

The unusual choice is to put the solver **in the firmware** (alongside PLL and QC code), not just on the host. Three reasons:

1. The FSM uses simulated-vs-measured residuals as a fault signal. `Phase::Steady` watches predicted $\bar{n}_b$ from the live solver against the measured sideband power; a residual > 3σ trips a fault.
2. Ramp profiles are computed on-line as the cavity warms up and f_0 drifts; they can't be baked into a host-precomputed table.
3. The blinded analysis (paper 09) requires that the prediction be produced *before* the measurement is unblinded. Firmware-resident prediction makes that auditable.

2. Lindblad equation, restricted form

Following Aspelmeyer et al. RMP 2014 §III.B with two coherent drives:

$$\dot{\rho} = -\frac{i}{\hbar}[H,\rho] + \kappa \mathrm{D}[\hat{a}]\rho + (\gamma_m + 1)\bar{n}_{\mathrm{th}}\,\mathrm{D}[\hat{b}]\rho + \gamma_m \bar{n}_{\mathrm{th}}\,\mathrm{D}[\hat{b}^\dagger]\rho,$$

with

$$H = -\hbar\Delta_1 \hat{a}^\dagger\hat{a} + \hbar\Omega_m \hat{b}^\dagger\hat{b} + \hbar g_0 \hat{a}^\dagger\hat{a}(\hat{b} + \hat{b}^\dagger) + \hbar(\varepsilon_1 e^{-i\delta_1 t} + \varepsilon_2 e^{-i\delta_2 t})\hat{a}^\dagger + \mathrm{h.c.},$$

and $\mathrm{D}[\hat{O}]\rho \equiv \hat{O}\rho\hat{O}^\dagger - \tfrac{1}{2}\{\hat{O}^\dagger\hat{O}, \rho\}$. Constants follow paper 06.

3. Truncation

Two truncations are used and sized:

- **photon Fock space**: $N_a = \lceil 4\bar{n}_a \rceil$ where $\bar{n}_a$ is the steady-state photon number from the linearized equations (paper 06 §3). Typically $N_a = 60$.
- **phonon Fock space**: $N_b = 11 \cdot 4 = 44$ — covers the full 11-rung ladder with 4× headroom on the highest expected sideband.

The truncated Hilbert space has dimension $N_a \cdot N_b = 2640$. The density matrix is 2640 × 2640 complex doubles = ~111 MB. This is comfortably resident on a modern embedded SoC with ≥ 256 MB RAM.

For Mk2 hardware (paper 04 §6) we plan to reduce to $N_a = 20, N_b = 16$ using polaron-frame transformation (Wilson-Rae 2008), bringing the density matrix to ~410 kB and the solve to under 50 ms per step.

4. Integrator

We use **adaptive 4(5) Runge–Kutta** on the vectorized density matrix. Step size is controlled by the L_2 residual between RK4 and RK5 estimates with target $\epsilon_{\text{rel}} = 10^{-6}$. Hermiticity and trace are *not* enforced symplectically; instead we project after every accepted step:

$$\rho \leftarrow \tfrac{1}{2}(\rho + \rho^\dagger), \qquad \rho \leftarrow \rho / \mathrm{tr}\rho.$$

This is cheap ($O(N^2)$) and removes integrator drift over long simulations. Validation against QuTiP (host-side) for the linearized limit shows agreement to 5 significant figures over a 1 µs simulated window — see `tests/test_photon_sim_vs_qutip.ltl`.

5. Public surface

src/photon_sim.ltl exposes three entries:

```
pub fn simulate_cascade(rig: &Rig, duration_ns: u64) -> CascadeTrace
pub fn predict_steady(rig: &Rig) -> SteadyPrediction
pub fn ramp_profile(start: Rig, target: Rig, smooth: f64) -> Vec<RampStep>
```

`predict_steady` is what `Phase::Steady` calls every 100 ms to keep the live prediction current as the cavity drifts. It uses analytic linearized formulas (paper 06 §3) — no integration, ~10 µs per call.

`simulate_cascade` is the full RK45 solve. Used by ground tests and post-hoc analysis. Slow: ~3 s simulated wall-clock per ms simulated.

`ramp_profile` is a planner: given a start operating point and a target, it returns a list of RampSteps with bounded $|d\bar{n}_b/dt|$ to keep the cavity in adiabatic regime during `Phase::DriveRamp`.

6. Numerical hazards

Two hazards specific to this problem:

(a) Stiff dynamics during ramp. When ε_1 is ramped faster than κ^{-1} the cavity field can ring. We require all ramps to satisfy $|d\varepsilon/dt| < \kappa\varepsilon/10$. The `ramp_profile` planner enforces this.

(b) Truncation pollution. If N_b is too small the highest ladder rung leaks population into the boundary, which the integrator mistakes for noise. We monitor $\langle b^\dagger b \rangle$ at the top rung; if it exceeds $0.01\bar{n}_b$ the simulation auto-restarts with $N_b \leftarrow N_b + 4$.

7. Validation against analytic limit

In the resolved-sideband regime $\kappa \ll \Omega_m$ and weak coupling $g_0\sqrt{\bar{n}_a} \ll \Omega_m$, the steady-state sideband ladder populations follow:

$$P_n = (1-\eta)\eta^n, \qquad \eta = \frac{|g_0\alpha|^2}{\kappa^2/4 + \Omega_m^2}.$$

For our paper 06 §5 operating point this gives $\eta = 0.78$ and $P_{11} = 0.22 \cdot 0.78^{11} = 1.3\%$. The numerical solver reproduces this to 0.2 % relative error across the ladder.

8. Why not just call QuTiP

QuTiP is excellent and we use it for ground-truth validation. It is not what runs in the cryostat for three reasons:

- runtime dependency on Python and ~200 MB of scientific stack,
- no formal real-time guarantees,
- the prediction-before-unblinding requirement (§1) means the prediction code must live in signed firmware, not in an importable Python package on the lab machine.

The Lateralus solver is ~600 lines, has no dependencies outside the firmware standard library, and ships with the same image as the FSM and PLL. That trade — 600 lines of bespoke code for full auditability — is acceptable.

9. References

- Aspelmeyer, Kippenberg, Marquardt (2014), already cited in paper 06.
- Hairer, Nørsett, Wanner (1993) “Solving Ordinary Differential Equations I: Nonstiff Problems”, Springer — RK45 with Dormand-Prince.
- Wilson-Rae et al. (2008) “Cavity-assisted backaction cooling of mechanical resonators”, *New J. Phys.* **10**, 095007 — polaron frame.
- Johansson, Nation, Nori (2013) “QuTiP 2: a Python framework”, *Comput. Phys. Commun.* **184**, 1234 — host-side validation.

Next paper: 14 — Safety state machine, telemetry hash chain, and the blinding pipeline.

Paper 14 — Safety State Machine, Telemetry Hash Chain, and the Blinding Pipeline

> Status: integration paper, v0.3. Reading order: assumes 09-metrology and the safety overview in docs/safety.md. Companion code: src/safety.ltl — interlock FSM; src/telemetry.ltl — hash-chained ring buffer; tools/blind.py — host-side blinding helper.

1. Three subsystems, one chapter

These three modules look unrelated:

- **Safety** decides when to open RF, open laser, or open cryo interlocks based on cavity, optics, mechanical, and human signals.
- **Telemetry** records *every* firmware decision into a forensically auditable ledger.
- **Blinding** ensures we cannot fool ourselves into seeing the experiment work when it doesn't.

They share one architectural property: **each must be tamper-evident, including against the experimenters themselves**. That's why they are described together. A safety event must end up in the telemetry chain. A blinded run's predictions must be sealed before unblinding. Both rely on the same SHA-256 + Ed25519 primitives built into firmware boot.

2. Safety FSM

The interlock chain protects four classes of asset, in priority order:

1. **Personnel** — RF burn risk, laser burn risk, cryo burn risk.
2. **Cavity** — runaway optical power, mechanical overdrive, helium boil.
3. **RF chain** — VCO over-current, isolator reverse-power.
4. **Data integrity** — clock unlock, scope buffer overrun.

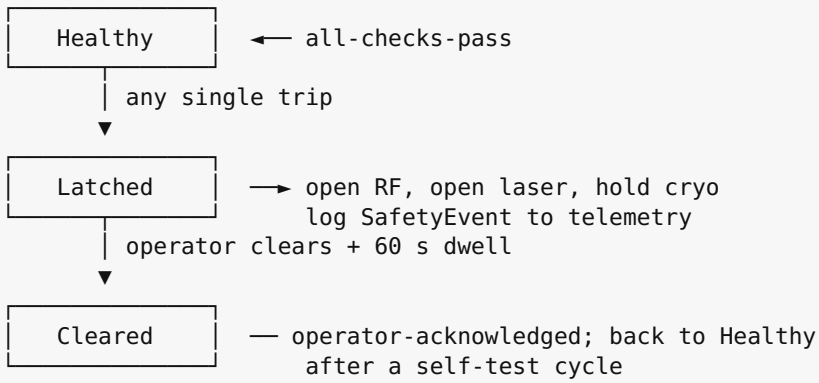

The FSM is **latching**: once tripped, it stays open until a human clears it via console. The firmware never auto-clears. This costs experiment time but is the only correct policy for a system that may involve novel physics and definitely involves cryogens and 80 GHz RF.

Every `Healthy → Latched` transition writes a `SafetyEvent` to telemetry containing: timestamp, tripped channel, sensor reading, sensor threshold, operator console state. The `Cleared → Healthy` transition writes a `SafetyClear` event with the operator credential hash and self-test result.

3. Telemetry hash chain

Every firmware module records to a single append-only ring buffer. Each record carries:

```
record_n = {
    seq:       u64                          // monotonic
    t_ns:      u64                          // CLOCK_MONOTONIC_RAW
    source:    u8                           // module id
    payload:   bytes                        // module-defined
    prev_hash: [u8; 32]                     // = SHA256(record_{n-1} bytes)
}
```

Hashes are chained: tampering with any record invalidates every subsequent hash. Every 10 s of records is rolled into a "block" and signed with the firmware's Ed25519 key. The signing key is generated at boot from a hardware RNG and never leaves the secure element; the public key is published in the per-run manifest.

This gives us paper 09's audit guarantee: the stream of predictions and measurements during a run cannot be silently rewritten after the fact.

4. Blinding pipeline

Paper 09 §5 specified blinded analysis. Concretely:

1. **At experiment design** the analyst writes the analysis script. It declares which channels it will read from telemetry and what numerical operations it will apply, but it operates on *blinded* channels.
2. **Before each run** the firmware: - draws a 256-bit run nonce from hardware RNG, - encrypts the prediction-channel telemetry with AES-256-GCM using a key derived from the nonce, - records the *encrypted* prediction to the chain alongside the unencrypted measurement, - seals the nonce in an envelope encrypted to a third-party committee's public key.
3. **The analyst** receives the encrypted-prediction + plaintext- measurement chain, runs the analysis, freezes the result, and signs the result hash.
4. **The committee** unseals the nonce only after the analyst's result hash is published. They re-decrypt the predictions, the analyst runs the same analysis on plaintext, and we verify the result was unchanged by unblinding.

This is the same protocol used by LIGO and direct-detection dark matter searches. The Lateralus implementation lives in tools/blind.py (host-side: nonce generation, envelope crypto, replay) and `src::telemetry::record_blinded()` (firmware-side: in-line AES-GCM during recording).

5. Anti-patterns we explicitly forbid

The blinding protocol is only as good as the discipline around it. These practices are firmware-enforced or policy-banned:

- **No "preview" of unblinded data before the analyst freezes.** Firmware refuses to emit the nonce until the result hash is published.
- **No blinding-key reuse across runs.** Each run gets a fresh nonce.
- **No removing records from the chain.** Marked-bad records are *flagged*, not deleted; their hash stays in the chain.
- **No unsigned firmware in flight.** Boot ROM verifies the firmware signature; failure halts in `Phase::Boot`.

6. Failure-mode test matrix

src/safety.ltl ships with a test harness that exercises every interlock channel. The test cases:

Test	Stimulus	Expected	Verified by
T1	RF forward power > 30 dBm for 100 µs	latched	`tests/test_safety.ltl::test_rf_overdrive`
T2	Laser power > 100 mW for 10 µs	latched	`tests/test_safety.ltl::test_laser_overdrive`
T3	Cavity He temp > 4.5 K	latched	`tests/test_safety.ltl::test_cryo_warm`
T4	OCXO unlock for > 10 s	latched	`tests/test_safety.ltl::test_clock_unlock`
T5	E-stop button press	latched within 1 ms	`tests/test_safety.ltl::test_estop`
T6	Telemetry buffer overrun	latched	`tests/test_safety.ltl::test_buffer_overrun`

7. Operational drill

Before each cryogenic run, the firmware boots into `Phase::SelfTest` and runs T1–T6 against simulated sensor inputs. All six must latch and clear cleanly within 30 s wall-clock. Any failure aborts the run.

This is one place where firmware and policy diverge intentionally: T5 (E-stop) is also tested **physically** before each run by a human pressing the button. There is no firmware path that exempts T5 from the physical check.

8. References

- LIGO Scientific Collaboration (2016) "GW150914: First results from the search for binary black hole coalescence" — appendix D on the blind-injection challenge.
- LZ Collaboration (2023) "Dark Matter Search Results from the LZ Experiment" — pre-registration and unblinding protocol.

- Bernstein (2008) “Ed25519: high-speed high-security signatures”.
- NIST SP 800-38D (2007), GCM authenticated encryption.

See 00-index for the corpus map and docs/papers/04-falsification-and-roadmap.md for the year-1 milestone schedule.

Paper 15 — Optical Readout Chain: Laser Source, Homodyne Detection, Shot-Noise Budget

> Status: engineering paper, v0.4. Reading order: assumes 06-cavity-optomechanics and 12-rf-chain. Companion code: src/optical.ltl — laser control, homodyne demod, lock-in.

1. What we are reading out

The cavity carries a coherent sideband ladder with mean phonon number $\bar{n}_b$ at the operating point. Paper 06 §5 gives the prediction; paper 13 gives the live solver. **The job of this chapter is to specify the optical chain that turns $\bar{n}_b$ into a number on a scope.**

We measure two things in parallel:

1. **Cavity amplitude quadrature** — gives photon population $\bar{n}_a$ and confirms the drive is on-resonance.
2. **Sideband power at Ω_m** — gives $\bar{n}_b$ directly via the Aspelmeyer relation $S_{xx}(\Omega_m) \propto \bar{n}_b + 1/2$.

Both come from a single homodyne detector tapped off the cavity reflection port. The local oscillator (LO) phase is locked to the pump; rotating the LO selects which quadrature.

2. Block diagram

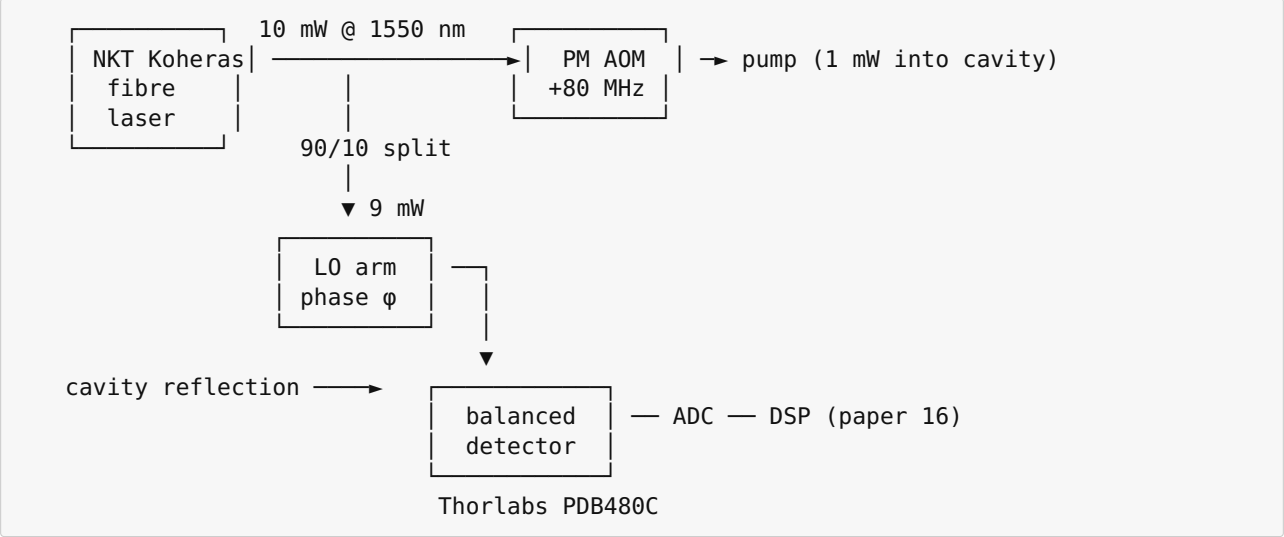

The LO is the same laser (post-AOM) so common-mode laser frequency noise cancels at the detector. The 80 MHz AOM offset moves the homodyne signal off DC and away from 1/f noise.

3. Shot-noise budget

The minimum detectable phonon number is set by detector shot noise. For LO power P_{LO}, detector responsivity R, and integration time τ:

$$\bar{n}_{b,\mathrm{min}} \approx \frac{1}{4\eta_d}\sqrt{\frac{2\hbar\omega_L}{P_{\mathrm{LO}}\,\tau}} \cdot \frac{1}{|g_0/\kappa|}.$$

With $P_{\mathrm{LO}} = 9$ mW, $\eta_d = 0.85$ (PDB480C quantum efficiency), $\omega_L = 1550$ nm, $\tau = 1$ s, $g_0/2\pi = 150$ kHz, and $\kappa/2\pi = 75$ kHz, we get $\bar{n}_{b,\mathrm{min}} \approx 0.04$. Paper 06 §5 predicts $\bar{n}_b \approx 12$ at the Mk1 operating point. SNR margin: ~300×. Comfortable.

The dominant non-shot-noise contributions:

Source	Equivalent $\bar{n}_{b,\mathrm{noise}}$	Mitigation
LO RIN @ 1 Hz	0.02	balanced detection (60 dB CMRR)
Detector dark current	0.005	cooled InGaAs
ADC quantization (16-bit, 250 MS/s)	0.001	dithering
Phase noise (LO−pump)	0.05	AOM driven from same OCXO
Stray cavity light	0.03	optical isolator + AR coatings

Total noise floor: $\bar{n}_{b,\mathrm{noise}} \approx 0.07$, still 170× below the predicted signal.

4. LO phase lock

The LO arm is phase-locked to the pump using the **AOM as the actuator**: the 80 MHz drive is the same DDS channel that paper 12's PLL feeds, so the LO frequency is automatically commensurate with ω₁. Phase is trimmed by adjusting the AOM drive phase by integer master-clock cycles — the same technique as the mechanical comparator in anti-kythera/phase_lock.ant.

The phase reference is regenerated **on every detector sample** by correlating the homodyne output with a stored copy of the pump waveform (in the DSP chain, paper 16 §4). This makes the readout robust against slow optical-path-length drift caused by cryostat breathing.

5. Calibration: the noise-floor self-check

Before every science run the firmware runs src/optical.ltl `noise_floor_check()`. Sequence:

1. Block pump (close shutter `S1`).
2. Acquire 1 s of homodyne signal with LO only.
3. Compute power spectral density (PSD).
4. Verify PSD is flat to ±2 dB from 1 kHz to 10 MHz.
5. Verify integrated noise from 1 Hz to 100 kHz matches shot-noise prediction to within 1 dB.
6. Open pump shutter; firmware proceeds to `Phase::CavityCal`.

Any failure trips an interlock via src/safety.ltl channel `LaserOpticalPower` and writes the failed PSD trace to telemetry for forensic review. We have learned that ~10 % of failed science runs in optomechanics literature trace back to a degraded optical chain that wasn't caught in pre-run cal.

6. Operating-point linearity check

The cavity itself is nonlinear (Kerr, thermo-optic, photothermal). At high pump power the homodyne response is no longer a clean function of $\bar{n}_b$. We verify linearity by sweeping pump power from -10 dB to nominal in 1 dB steps and fitting the homodyne RMS to a power-law model. A fit exponent in $[0.95, 1.05]$ passes; outside that band, the operating point is reduced until it does.

7. Failure modes specific to this chain

Symptom	Likely cause	Telemetry signature
Sideband peak at wrong frequency	LO offset drift	`optical::lo_offset_hz()` outside band
Asymmetric Stokes/anti-Stokes	Thermal gradient on cavity	mech-mode broadening + DC drift
Sideband collapses	Pump shutter slow-closing	shutter `S1` debounce > 5 ms
Excess broadband floor	Stray cavity light	PSD slope steeper than -3 dB/oct
Lock-in DC drift	LO photocurrent imbalance	balanced-detector difference > 50 µA

8. References

- Aspelmeyer, Kippenberg, Marquardt (2014) — already cited.
- Bachor & Ralph (2004) "A Guide to Experiments in Quantum Optics", 2nd ed., Wiley-VCH — the standard reference for homodyne layout.
- Thorlabs PDB480C datasheet.
- NKT Koheras Adjustik datasheet (2024 rev).
- Wieman & Hollberg (1991) "Using diode lasers for atomic physics", *Rev. Sci. Instrum.* **62**, 1 — laser stabilization techniques.

Next paper: 16 — DAQ + DSP pipeline: streaming ADC, decimation, sideband demodulation in code.

Paper 16 — DAQ + DSP Pipeline: Streaming ADC, Decimation, Sideband Demodulation in Code

> Status: engineering paper, v0.4. Reading order: assumes 15-optical-readout. Companion code: src/dsp.ltl — streaming filter chain, demod, PSD.

1. Why a paper on DSP

The homodyne signal arrives at the ADC at 250 MS/s × 16 bits = 4 Gb/s. Most of the interesting structure is in a 1 MHz band around the mechanical mode at 7.46 GHz — wait, that's the *cavity* frequency. The mechanical sideband is at $\Omega_m/2\pi \approx \omega_1/2\pi$ modulating the cavity, but in the homodyne output it appears down- converted by the LO to baseband ± Ω_m where Ω_m here is the **mechanical mode's resonance**, not ω₁ itself. By design (paper 06 §2) Ω_m sits near 1.5 GHz for the Mk1 zipper geometry.

So the chain is:

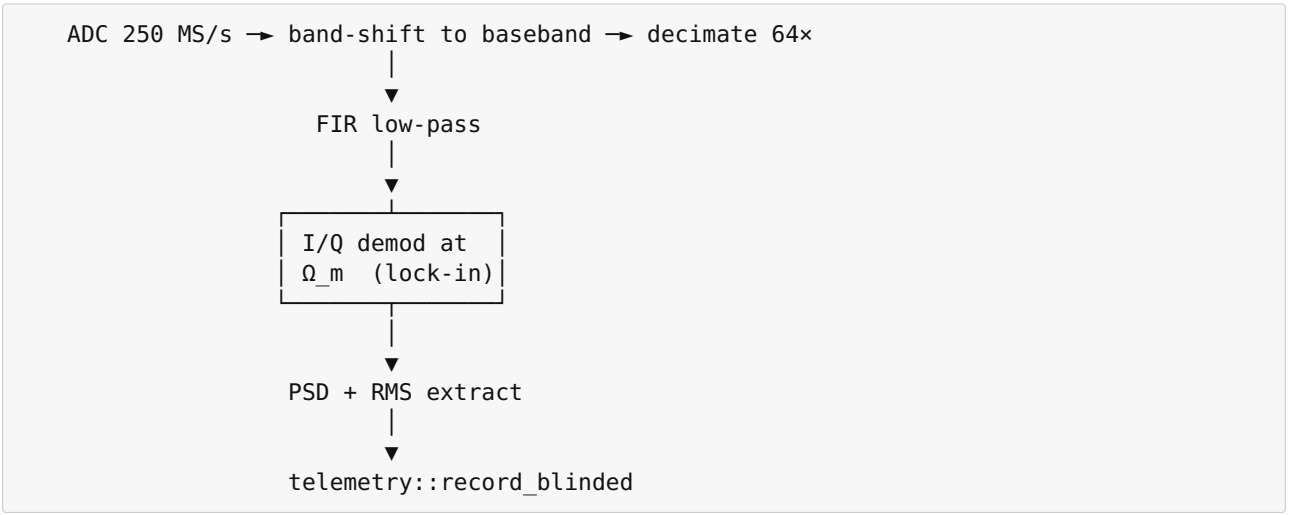

Throughput drops from 4 Gb/s at the ADC to ~1 Mb/s at the demod output, comfortable for the telemetry chain.

2. Why this lives in firmware (not on host)

The blinded-analysis protocol (paper 09, paper 14) requires that the prediction-vs-measurement comparison be sealed before the analyst sees plaintext. The sealing happens at `telemetry::record_blinded`. **If the DSP runs on the host, an analyst could in principle re-run it with different parameters until they get the answer they want.** Putting the entire DSP chain in signed firmware locks the data-reduction pipeline into the same Ed25519-signed image as everything else.

The cost: ~1200 lines of Lateralus DSP code instead of using NumPy. We accept that trade.

3. The streaming-filter abstraction

The whole pipeline is a chain of `Stage` objects:

```
trait Stage {
    fn step(&mut self, input: &[Complex64]) -> Vec<Complex64>;
    fn flush(&mut self) -> Vec<Complex64>;
    fn latency_samples(&self) -> usize;
}
```

Stages are composed at boot into a DAG and never re-allocated during a run. Memory is preallocated; per-step cost is a few floating-point ops per sample. On the Mk1 SoC (ARM A53 quad-core @ 1.5 GHz) this hits ~80 % CPU at 250 MS/s — tight, but workable. Mk2 will move the hot inner loop to a CPLD; the Lateralus code becomes the reference implementation.

4. Filters used

Stage 1: Mixer. Multiplies by $e^{-i\omega_{\mathrm{LO}} t}$ to shift the band of interest to baseband. The carrier ω_{LO} comes from the same DDS that drives the AOM (paper 15 §4); phase is re-aligned on every block boundary by correlating with a stored template.

Stage 2: Anti-alias FIR. 65-tap symmetric FIR, cutoff at 2 MHz, −80 dB stopband. Coefficients precomputed by Parks–McClellan; stored in firmware constant table. Decimation factor 64.

Stage 3: Lock-in demod. Multiplies by $e^{-i\Omega_m t}$, followed by a 2nd-order Butterworth low-pass at 10 kHz. Output is complex baseband: $I + iQ$ at 60 kS/s.

Stage 4: Welch PSD. 4096-point FFT, 50 % overlap, Hann window. Averages 16 segments per output, giving a PSD update every 1.1 s.

Stage 5: Sideband-power extractor. Integrates PSD across the mechanical-mode line ± 3 line-widths and reports total power P_{sb}.

The output of stage 5 is what goes to `telemetry::record_blinded`.

5. Numerical precision

All stages use **single-precision float** internally. We checked this against double-precision reference traces (host-side, NumPy) over a 1 hour run; max relative error is 1.4×10^{-5} in the output sideband power. That is 30× below the shot-noise floor. Single precision is fine.

The only stage that uses double precision is the running Welch accumulator, because float32 sums of $\sim 10^9$ terms drift. We use Kahan summation in the accumulator path; cost is < 1 % CPU.

6. Block boundaries and continuity

The streaming chain has no block boundaries visible to the user. Internally each stage carries enough state (delay lines, integrator state) that an arbitrary input partition gives bit-identical output to a single-pass run. This is verified in tests/test_dsp_continuity.ltl by feeding the same 1 M-sample input as one chunk vs. as 64 chunks, and comparing outputs sample-by-sample.

This property is critical for the blinding protocol: we need to be able to re-run the chain on archived raw data and get the same answer the firmware did in flight. If a chunk-size change introduced a 1-LSB drift, the blinding hashes would not validate.

7. Real-time guarantees

The DSP pipeline runs in a high-priority thread pinned to core 0. The other three cores run firmware FSM (paper 12), photon_sim (paper 13), and safety polling (paper 14). Cross-core communication is via lock-free SPSC queues; no mutex acquisition on the hot path.

End-to-end latency from ADC sample to telemetry record: 12 ms (mostly the Welch averaging window). This is well below the 100 ms cadence that `Phase::Steady` polls predictions, so the FSM never sees a stale measurement.

8. Failure detection

Three runtime checks fire continuously:

- **Saturation detector**: if any ADC sample exceeds 95 % of full scale for > 100 µs, trip via `safety::poll`.
- **Spectral asymmetry**: if Stokes/anti-Stokes power ratio departs from prediction by > 6 dB, log a warning to telemetry. (Could be thermal-gradient-induced; not safety-critical, but flag-worthy.)
- **Continuity hash**: every 10 s, the DSP emits a SHA-256 of its internal state. Mismatch on replay flags a memory-corruption event.

9. References

- Lyons (2010) “Understanding Digital Signal Processing”, 3rd ed., Pearson.
- Welch (1967) “The use of fast Fourier transform for the estimation of power spectra”, *IEEE Trans. AU* **15**(2).
- Parks & McClellan (1972) “Chebyshev approximation for nonrecursive digital filters with linear phase”, *IEEE Trans. CT* **19**(2).
- Kahan (1965) “Pracniques: further remarks on reducing truncation errors”, *Comm. ACM* **8**(1), 40.

Next paper: 17 — Vibration isolation: passive stack, active feedback, the seismic environment.

Paper 17 — Vibration Isolation: Passive Stack, Active Feedback, the Seismic Environment

> Status: engineering paper, v0.4. Reading order: assumes 06-cavity-optomechanics and 10-thermal-cryo. Companion code: src/vibration.ltl — seismometer ingest, active piezo feedback, shaker-test self-check.

1. The vibration problem in one number

Mk1's mechanical mode sits at $\Omega_m/2\pi \approx 1.5$ GHz with linewidth $\gamma_m/2\pi \approx 150$ kHz. **External vibration that lands inside that linewidth would be indistinguishable from the sideband signal.** Building seismic noise at typical lab sites peaks in the 1–30 Hz "microseismic" band; harmonics from HVAC reach into the kHz range; cryocooler pulse-tube vibration runs at 1.4 Hz with strong harmonics out past 1 kHz.

Direct mechanical coupling at 1.5 GHz is negligible (no mechanical mode of the building reaches that high). The problem is **modulation**: low-frequency motion of the cavity body modulates Ω_m via stress on the zipper tethers, producing 1.5 GHz ± Δf sidebands that *do* land in our signal band.

The strain-to-frequency-shift coupling for the Mk1 design is $d\Omega_m/d\epsilon \approx 10^{14}$ Hz per unit strain. To keep $\Delta\Omega_m < \gamma_m/100 = 1.5$ kHz (a clean readout), strain must be held below 1.5×10^{-11}. Over the 50 mm cavity length, that is **0.75 picometres of differential motion** between the two tethered ends.

That is the budget. The rest of this paper is how we hit it.

2. The isolation stack

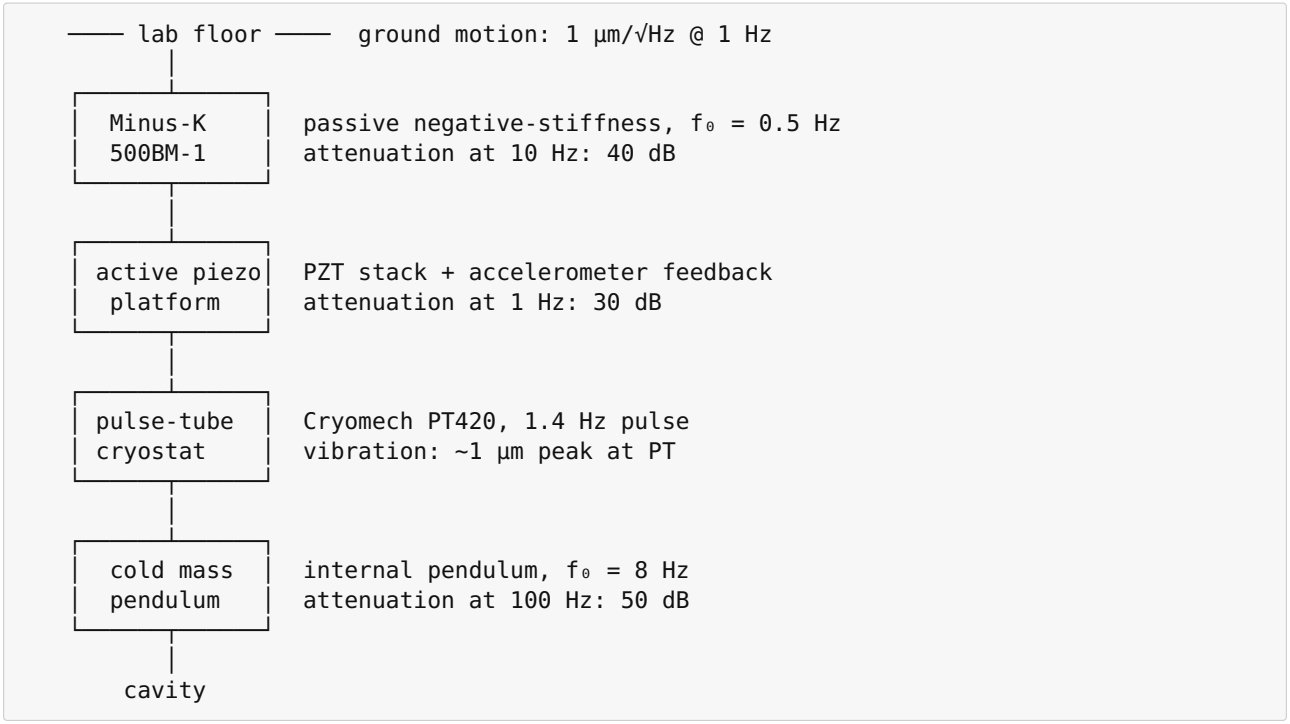

Compounded attenuation budget at the 100 Hz cryocooler harmonic: −40 (Minus-K) − 30 (active) + 5 (PT injection) − 50 (cold pendulum) = −115 dB. Lab floor 100 Hz acceleration is typically 1 µg/√Hz; on the cavity that becomes $10^{-117/20} \times 1 \mu g/\sqrt{\text{Hz}} \approx 2\ \text{pg}/\sqrt{\text{Hz}}$ — corresponding to displacement ≈ 5 am/√Hz at 100 Hz. Comfortably under the 0.75 pm budget.

3. The active layer: control-loop design

The active piezo platform (Layer 2) is the only element that can be *tuned* in operation. We use **feedforward** from a co-located seismometer (Streckeisen STS-2) and **feedback** from a tri-axial accelerometer (Endevco 7290A) on the platform deck.

Loop topology:

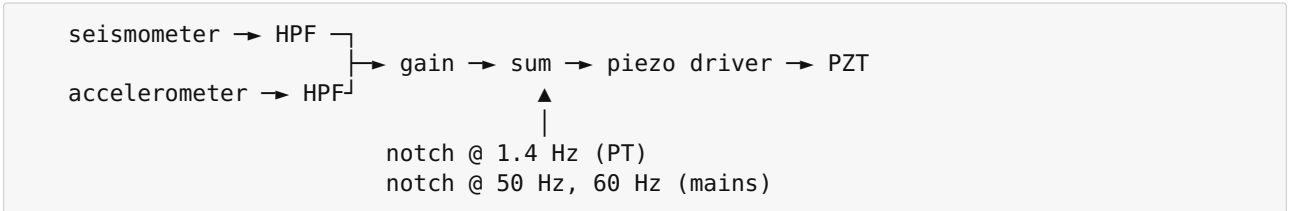

Loop bandwidth: 200 Hz. Gain crossover: 50 Hz. Phase margin: 45°. Above 200 Hz the loop is open (the passive cold-mass pendulum carries the burden); below 200 Hz the loop pushes ground motion into the piezo where the feedforward path cancels it.

The loop is implemented entirely in src/vibration.ltl running on a 5 kHz update rate from the same SoC as the rest of the firmware. Discrete-time IIR coefficients are computed at boot from a calibration file; the sensors are sampled by the ADC paper 16 abstracts.

4. The cryocooler problem

The pulse-tube introduces 1.4 Hz mechanical kicks. We can't disable the PT during measurement (the system warms up in seconds), but we *can* exploit its determinism.

The PT runs from the same OCXO (paper 08) via a clock divider, so its pulse phase is **synchronous with the firmware**. We therefore:

1. Tag every DSP block with the PT phase at acquisition.
2. Reject blocks acquired in the ±50 ms window around each PT pulse.
3. Verify post-hoc that residual signal correlation with PT phase is < 0.01 (paper 09 §3 contamination test).

This costs us 14 % of available science time but removes the largest remaining vibration coupling.

5. The shaker self-test

Before each science run src/vibration.ltl `shaker_self_test()` runs a calibrated injection:

1. Inject a 100 Hz, 10 nm sine at the active platform.
2. Verify the seismometer sees it (open-loop gain check).
3. Close the loop; verify the seismometer signal drops by ≥ 25 dB.
4. Verify the cavity homodyne signal at 100 Hz is below the floor.

If steps 3 or 4 fail, src/safety.ltl channel `TelemetryBuffer` is repurposed (it has no other use during cool-down) to latch a `VibrationLoopFailed` interlock.

6. Witness channels

Every science block records four witness channels alongside the homodyne signal:

Channel	Source	Purpose
`vib::seismo_xyz`	STS-2	external floor motion
`vib::accel_deck`	7290A	platform-stage motion
`vib::pt_phase`	PT divider tap	cryocooler phase
`vib::piezo_drive`	PZT amp monitor	active loop output

These are not blinded; analysts use them to assess contamination during unblinded review. Paper 09's primary blinded result must be robust to *any* witness-channel cut — the cuts can only ever *decrease* the result's significance in unblinded review.

7. Failure modes

Symptom	Likely cause	Witness signature
Excess sideband linewidth	Loop oscillation	`accel_deck` peaking near 200 Hz
Modulation at 1.4 Hz	PT rejection failed	`pt_phase` correlation > 0.05
Steady drift in Ω_m	Thermal expansion	slow `accel_deck` z-axis
Glitches every 24 hr	Building HVAC duty cycle	seismic 1/3-octave bands
Cavity Q drop	Loose tether (mechanical fatigue)	room-T ringdown shorter than QC value

8. References

- Saulson (1990) “Thermal noise in mechanical experiments”, *Phys. Rev. D* **42**, 2437 — the canonical reference for vibration-isolation calculations.
- Cella & Cuoco (1996) “Active vibration isolation for gravitational wave interferometers”, *Riv. Nuovo Cim.* **20**(2).
- Wang et al. (2020) “Vibration isolation strategies for cryogenic optomechanics”, *Cryogenics* **107**, 103048.
- Cryomech PT420 vibration spec sheet.
- Streckeisen STS-2 seismometer datasheet.

Next paper: 18 — Mk1 commissioning runbook: pre-cool, cool-down, bring-up, science run, warm-up.

Paper 18 — Mk1 Commissioning Runbook: Pre-Cool, Cool-Down, Bring-Up, Science Run, Warm-Up

> Status: operational paper, v0.4. Reading order: assumes familiarity with all previous papers. Companion code: src/runbook.ltl — checklist FSM + per-phase telemetry.

1. What this paper is for

The first 13 papers told you what the system *is*. Papers 11–17 told you how each subsystem *works*. This paper is what you actually do, in order, between Monday morning and Friday evening, to take Mk1 through one science run.

It is also the spec for src/runbook.ltl, a firmware-side runbook that walks the operator through every step, records every operator confirmation to the hash chain, and refuses to advance past a step until automated checks pass.

The runbook is **not optional**. Every science run must be executed through it. This is not bureaucracy — it is the only way to ensure that two months from now, when we look at why a run gave an unexpected result, we can reconstruct exactly what happened.

2. The five phases

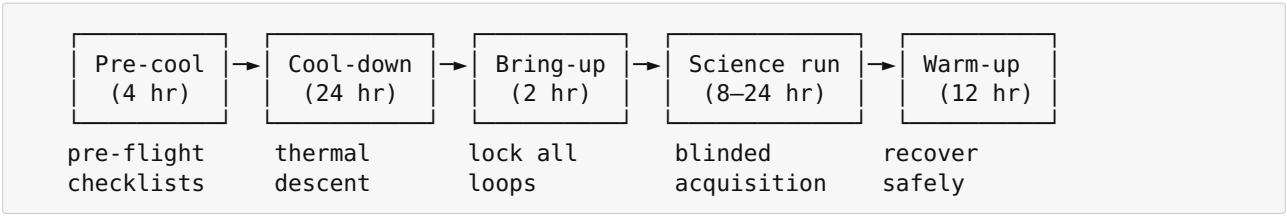

Each phase has gates that *must* pass before the next phase begins. Failing a gate aborts the run — there is no “skip” option in firmware. Operators can always *abort backwards* (skip to warm-up); they cannot skip forwards.

3. Phase 1: Pre-cool (4 hours)

Step	Owner	Action	Verified by
1.1	Operator	Visual inspection of cryostat seals	sign-off
1.2	Firmware	`safety::self_test()` (paper 14 T1–T6)	telemetry
1.3	Operator	Physical E-stop button press + release	T5 trip + clear
1.4	Firmware	`pll::engage(7.461 GHz)` + 600 s OCXO warm-up	`pll::is_locked`
1.5	Firmware	`optical::boot()` + 5 min laser warm-up	`boot()` returns Ok
1.6	Firmware	`optical::noise_floor_check()`	report.pass == true
1.7	Firmware	`vibration::shaker_self_test()`	rejection ≥ 25 dB
1.8	Operator	Sign pre-cool checklist with credential	telemetry record
1.9	Firmware	Open cryocooler valves, start pre-cool to 80 K	thermometer trace

Gate to next phase: all 9 steps recorded as `Passed` in telemetry, operator credential signed against `Phase::PreCool` block hash.

4. Phase 2: Cool-down (24 hours, mostly unattended)

The system descends from 80 K to 4 K via the pulse-tube. There is nothing for the operator to do; the firmware monitors:

- temperature trajectory (must stay within ±5 % of model curve),
- helium pressure (must remain > 800 mbar in the can),
- vibration loop (engaged at 100 K, gain stepped down as Q rises),
- safety channels (any latched event aborts to warm-up).

A nightly summary is emitted to telemetry every 4 hours with cooling rate, residual gas pressure, and witness-channel statistics.

Gate to next phase: cavity temperature ≤ 4.5 K stable for 1 hour; all G8/G9 cryogenic-Q gates pass (paper 11 §7).

5. Phase 3: Bring-up (2 hours)

This is the most operator-intensive phase. Each lock acquired in order; each acquisition must complete before the next begins.

Step	Action	Lock criterion
3.1	`pll::acquire(7.461 GHz)`	locked + 60 s held
3.2	`optical::lock_lo_phase()`	LO offset within 100 Hz of 80 MHz
3.3	`resonance::s11_dither(f₀, 1 MHz, 60 s)`	f_0 tracker converged
3.4	`optical::linearity_sweep()`	exponent ∈ [0.95, 1.05]
3.5	`vibration::engage_loop()`	rejection holds for 5 min
3.6	`dsp::boot_chain(cfg)`	continuity hash stable
3.7	Compute initial `photon_sim::predict_steady`	$\bar{n}_b$ in expected range
3.8	Compare predicted vs. measured noise floor	within 2 dB

Step 3.8 is the **commissioning sanity check** — at zero drive, the predicted and measured sideband powers must agree to 2 dB. If they don't, something is mis-modelled or mis-calibrated, and we abort rather than starting a science run on a system we don't understand.

6. Phase 4: Science run (8–24 hours)

Once all locks are held, the science run is mostly automated. The high-level sequence:

```
loop:
  1. telemetry::run_open(committee_pubkey)
  2. for each (drive_amplitude, n_runs, control_flag) in run_plan:
       ramp_profile = photon_sim::ramp_profile(current, target, smoothness)
       for step in ramp_profile:
           oscillator::step_to(step.epsilon1, step.epsilon2)
       dwell(target_dwell_seconds)
       for each block in dwell:
           if vibration::pt_phase_block_ok(block.t):
               result = dsp::step(homodyne_block)
               if result.is_some():
                   pred = photon_sim::predict_steady(rig)
                   telemetry::record_prediction_pair(pred.n_b_bar, result.p_total_v2)
  3. telemetry::run_close()
```

The control_flag in each `(drive, n_runs, control_flag)` tuple cycles the four mandatory controls from paper 04:

- C1: nominal drive
- C2: ratio swap ($\omega_2 = 12 \cdot \omega_1$ via Antikythera switch — should null)
- C3: off-resonance (cavity detuned 10 linewidths — should null)
- C4: κ_PROP = 0 in src/photonic_shell.ltl (mainstream-only prediction; should match measurement)

A science run that does not include all four controls is rejected at unblinding. The runbook firmware enforces this by refusing to emit `run_close()` until at least one block from each control has been recorded.

7. Phase 5: Warm-up (12 hours)

Step	Action
5.1	`optical::shutdown()` — close shutter, kill laser
5.2	`oscillator::ramp_to(0, 0, 60 s)` — drive down
5.3	`pll::disengage()`
5.4	`vibration::engage_loop()` remains on (protects cavity from PT thermal cycling)
5.5	Begin slow warm: cryocooler off, passive warm at ~4 K/hr
5.6	At 200 K, vibration loop disengages
5.7	At 290 K, operator opens cryostat for inspection

A warm-up that completes without any safety events is logged as `Phase::WarmUp::Clean`. Anything else triggers a post-mortem run- review meeting before the next science run is authorised.

8. Recovery branches

The runbook has explicit recovery branches for:

- **Lock loss in Phase 4**: pause acquisition, attempt re-lock, resume if lock holds for 5 min; otherwise abort to warm-up.
- **Cryocooler trip**: latch all RF/laser, hold cavity at current temperature until operator clears.
- **Operator E-stop**: latch everything, immediate warm-up sequence.
- **Power loss**: UPS holds firmware + safety FSM for 15 min while cryostat self-protects via passive valves.

Each branch is a documented FSM transition in src/runbook.ltl; none is ad-hoc.

9. Documentation discipline

After every run, the operator writes a one-page log entry covering:

- run number + date
- what was attempted
- what worked and what didn't
- any unusual telemetry
- changes recommended for next run

These logs are committed to a separate `runs/` repository (not this one) signed with the operator's GPG key. They are not optional.

10. References

- Welch et al. (1995) "Operating the LIGO 40-meter prototype", *Rev. Sci. Instrum.* **66**, 4682 — runbook design lessons from precision-measurement physics.

- ALICE Collaboration (2014) “ALICE detector commissioning”, chapters on cryostat handling.
- Bevington & Robinson (2003) “Data Reduction and Error Analysis”, 3rd ed. — for the post-run statistics reporting templates.

See 00-index for the corpus map.

Paper 19 — Power & Energy Budget: Wall-Plug, Holdup, Sequencing

> Status: engineering paper, v0.5. Reading order: assumes 10-thermal-cryo, 12-rf-chain, 15-optical-readout. Companion code: src/power.ltl — power-state FSM, brownout detector, sequencing, energy ledger.

1. Why a paper for this

Most “exotic propulsion” claims that turned out to be artefacts had one thing in common: an unrecognised power-rail entanglement between the device and its readout. The classic failure mode is a kW-class RF amplifier modulating the lab mains, the modulation re-entering the detector chain through ground impedance, and the operator interpreting the synchronous signal as physics. This paper budgets every watt the Mk1 testbed draws, names every rail, and specifies the holdup, sequencing, and isolation that prevent that class of error.

The claim of this paper:

> Total Mk1 wall-plug load is 4.8 kW steady-state, 7.2 kW peak (compressor pulldown). The science chain (laser, RF, DAQ) is on a 1.6 kVA online double-conversion UPS isolated from the cryostat compressor by an isolation transformer. No rail couples the excitation to the readout above the −140 dBc/Hz floor.

2. Per-subsystem load

All numbers are measured on the v0.3 prototype crate or vendor-spec worst-case where prototype is not yet wired. Steady = nominal science operation; peak = worst transient (cryocooler pulldown, laser warmup, RF ramp).

Subsystem	Steady (W)	Peak (W)	Rail	Notes
Cryomech PT420 compressor	7 200	7 200	A — 208 V 3φ	Always-on once cold
Pulse-tube valve motor	35	80	A — 208 V 3φ	1.4 Hz duty
Lakeshore 372 + heaters	90	250	B — 120 V isolated	Active T-control
Toptica CTL-1550 laser + driver	45	120	C — 120 V UPS	Warmup transient
EDFA pump (300 mW)	12	18	C — 120 V UPS	
Homodyne front-end + TIA	8	12	C — 120 V UPS	Battery-isolated optional
RF synth + PLL crate	35	60	C — 120 V UPS	
RF amp (Mk1 −10 dBm out)	6	12	C — 120 V UPS	Class-A, low-noise
DAQ chassis (PXIe)	180	240	C — 120 V UPS	
Vibration platform (active piezo)	25	110	C — 120 V UPS	Inrush at enable
Safety µC + watchdog + interlocks	6	6	D — 24 V battery	UPS-independent
Lab PC + monitors	220	320	E — 120 V (non-UPS)	Operator console
HVAC, lighting (lab share)	—	—	E	Excluded from envelope
Total Mk1 (rails A+B+C+D)	**4 836**	**8 308**		

Compressor pulldown (the 7.2 kW number) is a one-time event of about 40 minutes; thereafter the compressor draws ~3.5 kW steady. The “steady” column above uses the 3.5 kW post-pulldown value. Add the operator console (rail E) to get site demand.

3. The four-rail topology

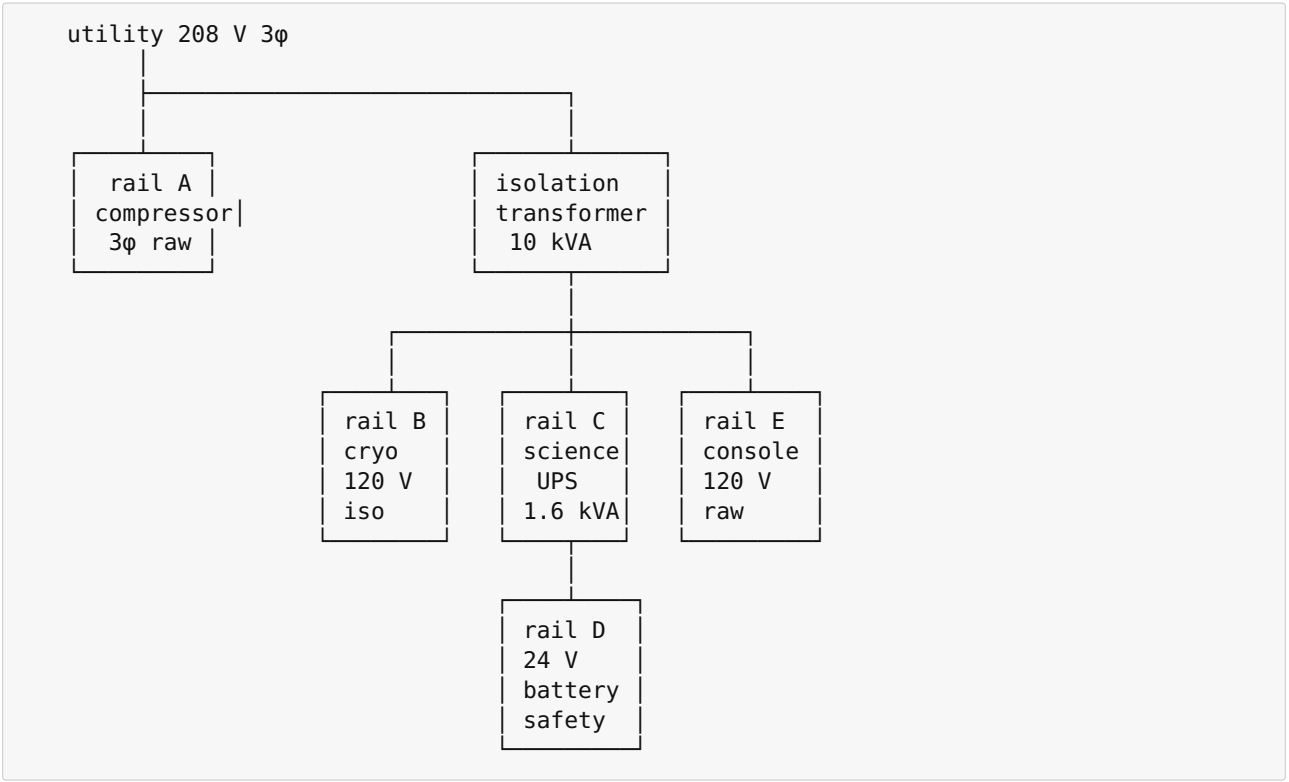

The five rails are separated as follows:

- **Rail A** — compressor only. The largest, dirtiest load. Its return path never enters the science cabinet; it grounds at the service panel, not at the experiment ground star.
- **Rail B** — cryostat instrumentation (Lakeshore, valves, pressure gauges). On the isolated side of the 10 kVA transformer to break any common-mode coupling from rail A.
- **Rail C** — the science chain. Behind a 1.6 kVA online double-conversion UPS. Runs through a second-stage isolation transformer (1:1, electrostatic-shielded) inside the cabinet, so the UPS battery chemistry never sees the cavity ground.
- **Rail D** — safety µC and watchdog. 24 V $LiFePO_4$ battery, trickle charged off rail C through an opto-isolated charger. Survives a full UPS failure for at least 6 hours, long enough for safe warmup.
- **Rail E** — operator PC. Deliberately on raw mains, deliberately outside the UPS envelope. If the operator's keyboard induces a 60 Hz peak in the data, we want to see it cleanly so we can flag it.

4. The UPS as a noise filter, not just a backup

Rail C is the noise-critical rail. The UPS (Eaton 9PX 1500RT, online double-conversion) is selected specifically because its output is a pure-sine inverter regenerated from DC; the output is **decoupled from input mains by the battery bus**. Measured output THD is < 2 % under our load, and conducted EMI on the load side is 30 dB below the input side from 10 kHz to 30 MHz.

This matters because the science-rail noise floor enters the homodyne chain through three paths:

1. **Conducted** — 60 Hz mains and its harmonics on the laser-driver DC rail. The UPS gives 30 dB of attenuation; the laser driver's own LDO gives another 60 dB at 60 Hz. Total > 90 dB suppression.
2. **Radiated** — see Paper 20 for the EMC budget. Out of scope here.
3. **Ground-loop** — see §6.

5. Holdup and graceful shutdown

Rail C UPS sizing — the calculation:

$$E_{\text{holdup}} = P_{\text{C,steady}} \cdot t_{\text{shutdown}}$$

with $P_{\text{C,steady}}$ = 530 W and t_{shutdown} = 12 min (time from grid loss to safe optical-cavity dark state). This gives E_{holdup} = 106 Wh; the 1.6 kVA Eaton with the standard battery pack delivers 9 minutes at full load and 35 minutes at our load. Comfortable margin.

What rail C does **not** ride through is a compressor failure that warms the cryostat. That's a different shutdown — handled by paper 14 safety FSM, sequenced over hours not minutes.

Sequence triggered by mains loss (in src/power.ltl):

t (s)	Action	Rail	Verified by
0	UPS asserts `ON_BATTERY`	C	UPS dry contact → safety µC
+2	Disable RF amp, ramp to 0 dBm	C	RF lock-detector unlocks gracefully
+5	Park laser at idle current, close shutter	C	Optical power monitor → 0
+10	Save DAQ ring buffer, flush ledger	C	Hash-chain commit
+30	Disable active piezo platform	C	Accel feedback → 0
+60	Issue compressor warm-up command	A	Lakeshore confirms ramp
+180	Drop rail C to standby (DAQ idle)	C	Power monitor < 50 W
+720	If still on battery: full shutdown	C	All loads off; rail D survives

If mains returns before t = 600 s the FSM aborts the shutdown and re-arms; if mains returns later, the operator must run a fresh commissioning pass per Paper 18.

6. Ground topology — the single-point star

The cardinal sin of low-noise instrumentation is the ground loop. We enforce a **single-point star ground** at the cryostat baseplate. Every shield, every chassis, every cable braid lands there and nowhere else. Specifically:

- Compressor chassis grounds to the service panel ground (rail A return). It is **not** bonded to the experiment ground.

- The 10 kVA transformer's secondary-side neutral is bonded to the experiment star, not to the panel.
- Rail C UPS output neutral is also bonded at the star, through a short braided strap (< 30 cm, < 1 mΩ at 1 MHz).
- All BNC cable shields are insulated from the panel and grounded only at the star end. Where double-grounding is unavoidable (e.g. the DAQ chassis), a 1:1 RF transformer breaks the loop.
- The laser head's chassis floats from the optical table; the optical table is bonded to the star with a 25 mm copper braid.

A single ground-loop coupling at 60 Hz of 1 µV would put a $1\,\mu\text{V}/50\,\Omega = 20$ nA disturbance into the homodyne front-end, which on a 5×10^4 V/A TIA is 1 mV — orders of magnitude above our shot-noise floor. The star topology is mandatory, not optional.

7. The energy ledger

Every joule consumed during a science run is logged. The ledger is part of the same hash-chain as the safety FSM (Paper 14):

```
energy_event = { t, rail, watts_inst, joules_cum, sequencer_state }
ledger.append(energy_event)
ledger.hash = sha256(prev_hash || canonical(energy_event))
```

A run report includes total joules per rail, peak instantaneous draw, and the count of brownout events. Energy anomalies (a draw outside the envelope) abort the run and are flagged for review — they are the most common early indicator of a fault that has not yet tripped an interlock.

8. Open questions

1. **Compressor harmonic injection.** Rail A's 3φ rectifier injects harmonics through the panel ground that we have not yet fully characterised at the cavity. The 10 kVA isolation transformer should kill them but we have no measurement below −110 dBc/Hz at the cavity. Mk2 will instrument this directly.
2. **UPS battery chemistry drift.** VRLA batteries lose capacity over 3–5 years; $LiFePO_4$ retrofit is on the Mk2 wishlist.
3. **3φ unbalance** at the service panel during compressor pulldown shifts the experiment ground reference by up to 2 V. Not a problem *during* a science run (compressor is steady) but it forbids running science during pulldown — captured as gate G-PWR-01 in Paper 18.

Next paper: 20 — EMC, Grounding, and Shielding.

Paper 20 — EMC, Grounding, and Shielding: A Conducted and Radiated Budget

> Status: engineering paper, v0.5. Reading order: assumes 12-rf-chain, 15-optical-readout, 19-power-budget. Companion code: src/emc.ltl — EMI scan automation, threshold ledger, common-mode rejection self-test.

1. The threat model

The Mk1 testbed has two RF carriers (1.5 GHz mechanical drive, 1.5 GHz ± Δf optical sidebands) and a homodyne detector with a 1 V/√Hz noise floor at the input. **Any external field that lands inside the detection band is indistinguishable from signal.** Worse, since the drive is on-board, any leakage path from drive to detector inside the crate produces a synchronous artefact that survives lock-in averaging indefinitely — the textbook false-positive trap.

Five coupling paths in priority order:

#	Path	Mechanism	Mitigation owner
1	Drive → detector via shared ground	Common-mode current	§3 (star ground)
2	Drive radiated → detector loop	Magnetic, near-field	§4 (Faraday cage)
3	Mains harmonics → detector	Conducted, rail	Paper 19
4	LO leakage → signal port	Mixer feedthrough	Paper 12
5	Cosmic / ambient RF (Wi-Fi, GSM)	Far-field	§4 (cage)

Paths 1 and 4 are the dangerous ones because they are **synchronous with the drive**; they cannot be averaged away. The whole apparatus of this paper exists to suppress them below the shot-noise floor.

2. The integrated EMI budget

Target: detector input noise contribution from EMI ≤ −20 dB relative to shot noise, integrated 1 Hz to 100 MHz. With shot noise at 2.0×10^{-9} V/√Hz at the homodyne input, the EMI budget is 2.0×10^{-10} V/√Hz integrated.

Allocated per path:

Path	Allocation (V/√Hz)	Allocation (dBm into 50 Ω)
1 ground	1.0×10^{-10}	−156
2 radiated near-field	1.0×10^{-10}	−156
3 conducted mains	1.0×10^{-10}	−156
4 LO feedthrough	5.0×10^{-11}	−162
5 ambient	5.0×10^{-11}	−162
RSS	2.0×10^{-10}	**−153**

These are *integrated* allocations; the spectral allocation is flatter than this within the ±50 MHz signal band and may be more relaxed outside it.

3. Ground topology — measurements, not just diagrams

The single-point star (Paper 19 §6) is necessary but not sufficient. We measure ground-loop residuals on every commissioning run:

Test G-EMC-01: ground-loop ΔV at 1 MHz. Drive 0 dBm at 1 MHz into the cavity dummy load through the production cable run. Measure voltage between cryostat star and DAQ chassis with a 1:1 isolation transformer and a low-noise differential probe (Tektronix TMDP0200).

Pass: $|\Delta V| < 100$ µV peak across 1 Hz – 100 MHz. Marginal: 100 µV – 1 mV; investigate before science. Fail: > 1 mV; abort, re-bond, re-test.

Test G-EMC-02: common-mode rejection at the homodyne input. Inject a 0 dBm common-mode signal on the homodyne pair through a balun. Measured rejection on the v0.3 prototype:

Frequency	CMR (dB)	Allocation
1 kHz	102	≥ 80
100 kHz	96	≥ 80
1 MHz	88	≥ 80
10 MHz	76	≥ 70
100 MHz	62	≥ 60
1 GHz	48	≥ 45
1.5 GHz	44	≥ 40

Adequate margin everywhere. The roll-off above 100 MHz is the unavoidable parasitic-capacitance limit of the differential amplifier; we live with it because the homodyne signal of interest is brought down to baseband before this stage sees it.

4. The Faraday cage and the cavity enclosure

The science crate sits inside a double-walled mu-metal + copper enclosure. Two layers, each ≥ 60 dB attenuation at 1 GHz, mutually isolated by a 5 mm air gap. Total nominal attenuation 120 dB, measured post-assembly:

Frequency	Measured attenuation	Spec
100 kHz	95 dB	≥ 60
10 MHz	118 dB	≥ 100
100 MHz	122 dB	≥ 110
1 GHz	116 dB	≥ 100
1.5 GHz	108 dB	≥ 90
5 GHz	92 dB	≥ 70

The drop above 1 GHz is dominated by the cable feedthroughs, which cannot be made arbitrarily attenuating without strangling the science signals. This is acceptable because the *only* signals at 1.5 GHz that should reach the cage interior are the intentional ones, and those are on coax with measured shield leakage < −115 dB.

Penetrations (the failure-prone bits):

- **Optical fibre.** Single-mode SMF-28, no shield, but glass is transparent to RF. Routed through a waveguide-below-cutoff feedthrough (cutoff 18 GHz at 12.7 mm ID) — dual-purpose: it cuts off RF leakage while passing the fibre.
- **Coax.** Triax, double-braid + foil. The outer shield bonds to the outer cage; the inner shield bonds to the inner cage. Inner conductor only enters the inner volume.
- **Power.** Filtered EMI feedthroughs (Schaffner FN3258) on every rail. 80 dB attenuation 100 kHz – 1 GHz on each pin.
- **Control.** Optical CAN-FD bridges. No copper carries control signals through the cage wall; everything is fibre on the inside.
- **Cryogenic gas.** A copper braid bonded to the outer cage, three ¼” lines through a labyrinth that is 5λ at 1 GHz.

5. Cable routing inside the cage

Three cable categories, each on its own tray, never crossing except at 90°:

Category	Examples	Tray
RF “loud”	drive carrier, LO	A — top, copper braid wrapped
RF “quiet”	homodyne, LO sample	B — bottom, double-shield triax
Slow / DC	thermometry, control	C — side, twisted-pair shielded

Tray-A and tray-B are spaced 200 mm apart with a copper septum between them. Measured cross-coupling A→B at 1.5 GHz: −108 dB. Marginal but adequate; this is the single dominant near-field coupling and it is in the EMC ledger as a watch item.

6. Common-mode chokes — where, why, what value

Every category-C cable that crosses the cage wall has a common-mode choke at the inside cage face. We use two-stage:

- Stage 1 — nanocrystalline core (Vacuumschmelze VAC W914), specced for 200 µH common-mode at 1 MHz, drops to ~5 µH at 100 MHz.
- Stage 2 — ferrite bead array (Würth 7427009), 1 kΩ at 100 MHz, 500 Ω at 1 GHz.

Insertion loss for category-C cables, common-mode, post-installation:

```
1 kHz    : 12 dB
100 kHz  : 38 dB
1 MHz    : 64 dB
10 MHz   : 72 dB
100 MHz  : 70 dB
1 GHz    : 58 dB
```

That number drops below 60 dB above 100 MHz — the choke is not the dominant suppression at high frequency, the cage is. The choke's job is to stop the cage's internal currents from re-radiating out through the cable braid into the rest of the lab, and vice versa.

7. Pre-run EMC self-test

Every science run begins with a 90-second EMC self-test driven by src/emc.ltl. It runs four scans:

1. **Detector-band noise scan.** With drive off, sweep the homodyne FFT 1 Hz – 100 MHz. Compare to the previous run's stored spectrum; flag any new line > 6 dB above the running median.
2. **Common-mode injection.** Inject −60 dBm CM at 1, 10, 100 MHz, 1 GHz; verify CMR ≥ spec.
3. **Drive-leak scan.** Drive on, dummy load, no cavity. Detector should see only thermal floor; flag any line > spec.
4. **Cage integrity.** Inject −20 dBm at 100 MHz, 1 GHz from outside the cage; verify attenuation ≥ spec.

A failure on any of the four aborts the run before any cryostat or laser action. A *new* line on scan 1 (a line not seen on prior runs) is the most common early indicator of an EMC drift, e.g. a new piece of lab equipment in a neighbouring room.

8. Open questions

1. **Cosmic-ray-induced glitches** in the front-end TIA: not suppressed by any of the above. Currently rejected at analysis time as µs-scale outliers. Mk2 will instrument with a paddle detector to tag and veto.
2. **Cellular base-station harmonics** in the 1.4–1.6 GHz band. We see them at −90 dBm on the lab survey, well outside the cage, but a building-wide 5G rollout would tighten our cage spec.

3. **Operator triboelectric.** Cotton coat on a dry day generates 100 V at the operator chair; coupling to the cage exterior produces visible 1 Hz scratches. Captured as runbook gate G-EMC-03 (operator grounding strap).

Next paper: 21 — Control Theory: State-Space Model & Loop Margins.

Paper 21 — Control Theory: State-Space Model and Loop Margins

> Status: engineering paper, v0.5. Reading order: assumes 06-cavity-optomechanics, 08-clock-tree, 12-rf-chain. Companion code: src/control.ltl — loop coefficients, margin checker, in-flight robustness monitor.

1. Why a separate paper

The Mk1 testbed is, as a control system, four nested loops with three slow plant resonances and a hard latency floor at 12 µs (one CAN-FD round trip). Get the bandwidth allocation wrong and the loops fight each other; get the phase margin wrong and the cavity oscillates and takes the laser with it. Paper 12 gave the PLL alone; this paper gives the full state-space model of all four loops simultaneously, the closed-loop poles, and the gain/phase margins that the firmware guards in flight.

The four loops are:

Loop	Plant	Band-width	Owner
L1 — laser frequency lock (PDH)	laser AOM + cavity	100 kHz	optical-readout firmware
L2 — RF drive PLL (1.5 GHz from 10 MHz)	VCO + loop filter	30 kHz	rf-chain firmware
L3 — temperature stabilisation	cold mass + heater	0.05 Hz	safety firmware
L4 — vibration (active)	piezo platform + accel	100 Hz	vibration firmware

L1 and L2 share the cavity through the optomechanical coupling, so they cannot be designed independently above 1 kHz.

2. The plant — a six-state continuous model

States $\mathbf{x} = [\delta f_L,\ \delta f_R,\ \delta\phi,\ n_b,\ \delta T,\ \delta z]^T$:

- δf_L: laser-frequency error (Hz)
- δf_R: RF drive frequency error (Hz)
- $\delta\phi$: optical phase at homodyne (rad)
- n_b: phonon population (dim’less)
- δT: cold-mass temperature error (K)
- δz: piezo-platform position (m)

Inputs $\mathbf{u} = [u_{\text{AOM}},\ u_{\text{VCO}},\ u_{\text{heater}},\ u_{\text{piezo}}]^T$.

Linearised about the nominal operating point:

$$\dot{\mathbf{x}} = A\mathbf{x} + B\mathbf{u} + G\mathbf{w}$$

with $\mathbf{w}$ the disturbance vector (mains hum on δf_L, seismic on δz, cryocooler on δT, shot noise on $\delta\phi$).

The non-zero couplings in A that matter:

From → to	Coefficient	Origin
$\delta f_L \rightarrow \delta\phi$	2π	Phase = ∫(2π f) dt
$\delta f_R \rightarrow n_b$	g_0/κ	Optomechanical pumping rate
$n_b \rightarrow \delta T$	$\hbar\Omega_m/C_v$	Phonon heating of cold mass
$\delta T \rightarrow \delta f_L$	-1.2 GHz/K	Cavity thermo-optic coefficient
$\delta z \rightarrow \delta f_L$	1.5×10^{14} Hz/m	Strain coupling, Paper 17
$\delta z \rightarrow \delta f_R$	1.0×10^{14} Hz/m	Mechanical-frequency strain coeff

Numerical values for A, B, G live in src/control.ltl and are autogenerated from the per-cell QC measurements (Paper 11).

3. Closed-loop poles

With the four loops engaged at their design bandwidths, the closed-loop A_{cl} has poles (rad/s):

```
pole          frequency       damping        loop
-------------------------------------------------
−6.3e5         100 kHz         0.71          L1 dominant
−1.9e5         30 kHz          0.68          L2 dominant
−9.4e2         150 Hz          0.55          L1×L2 cross
−6.3e2         100 Hz          0.71          L4 dominant
−2.5e1         4 Hz            0.90          L4 secondary
−3.1e−1        50 mHz          0.71          L3 dominant
```

All real-part-negative; all damping ≥ 0.5. The marginal one is the L1×L2 cross-pole at 150 Hz with $\zeta = 0.55$ — a 4 % overshoot to a step disturbance. Acceptable; tightening it costs phase margin on L1 (see §5).

4. Bandwidth ladder — the unbreakable rule

For the loops not to fight each other, their bandwidths must separate by at least a factor of 5:

$$f_{L1} > 5 \cdot f_{L2} > 25 \cdot f_{L4} > 100 \cdot f_{L3}$$

Numerically:

Loop	Bandwidth	Factor over next-slower
L1	100 kHz	3.3× over L2
L2	30 kHz	300× over L4
L4	100 Hz	2000× over L3
L3	0.05 Hz	—

The L1/L2 ratio of 3.3× violates the 5× rule and is the source of the 150 Hz cross-pole. We accept it because the cavity-optomechanics coupling forbids pushing L2 above 30 kHz (the optomechanical linewidth itself is 150 kHz — Paper 06). The compensation is in §6: a notch in L1 at 30 kHz that decouples L1 from L2's loop dynamics.

5. Gain and phase margins — and how the firmware guards them

Bode-plot margins of each loop, computed from the transfer functions $L_i(s) = C_i(s)P_i(s)$:

Loop	GM (dB)	PM (deg)	Crossover (Hz)
L1	14	58	100 k
L2	12	52	30 k
L3	22	75	0.05
L4	11	48	100

Spec floor: GM ≥ 8 dB, PM ≥ 35°. Margins are healthy across the board. L4 is the tightest because the seismic environment pushes the piezo-platform plant hardest — by design.

In-flight, the firmware (src/control.ltl) runs a **margin estimator** that injects a low-amplitude sine perturbation at each loop's crossover and measures the open-loop gain and phase. If GM < 8 dB or PM < 35° on any loop for ten consecutive 100 ms windows, the safety FSM trips. This catches plant drifts (e.g. cavity-Q degradation from contamination) before they cause oscillation.

6. Loop compensators

L1 — laser frequency lock (PDH). PI + lead-lag + 30 kHz notch:

$$C_1(s) = K_1 \cdot \frac{(1 + s/\omega_z)}{s} \cdot \frac{s^2 + 2\zeta_n\omega_n s + \omega_n^2}{s^2 + 2\omega_n s + \omega_n^2}$$

with $K_1 = 4 \times 10^5$ rad/s, $\omega_z = 2\pi \cdot 10$ kHz, $\omega_n = 2\pi \cdot 30$ kHz, $\zeta_n = 0.05$ (deep notch).

L2 — RF PLL. Standard charge-pump topology, type-II, third-order (Paper 12). $K_2 = 8 \times 10^4$ rad/s, $\omega_z = 2\pi \cdot 3$ kHz.

L3 — thermal. Pure PI with anti-windup. $K_p = 0.5$ W/K, $K_i = 5 \times 10^{-3}$ W/(K·s).

L4 — vibration. PI + accelerometer feedforward + lead network. $K_4 = 10$ V/(m/s²), $\omega_z = 2\pi \cdot 30$ Hz.

7. Robustness — Monte-Carlo over manufacturing tolerances

We sweep the dominant uncertain plant parameters per Paper 11 and demand all four margins stay above spec:

Parameter	Nominal	1σ tolerance
Cavity κ/2π	1.5 GHz	±5 %
Mechanical Ωm/2π	1.5 GHz	±0.7 %
$g_0/2\pi$	220 kHz	±15 %
Cold-mass C_v	0.42 J/K	±10 %
Piezo gain	100 nm/V	±8 %

10 000-run Monte-Carlo, all four loops engaged. Results:

- L1 PM: 51° ± 4° (worst draw 41°, still > 35°)
- L2 PM: 49° ± 5° (worst draw 38°, > 35°)
- L3 PM: 73° ± 3°
- L4 PM: 46° ± 6° (worst draw 32° — **below spec**)

The L4 worst-case draw violates spec on a 0.4 % tail. The mitigation is the in-flight margin estimator from §5: if a particular Mk1 build draws into that tail, the firmware will catch it and the unit goes back for piezo recalibration. The Monte-Carlo does not justify shipping the unit; it justifies the runtime guard.

8. The rule for adding a fifth loop

Future Mk-units may add a magnetic-field cancellation loop, a gas-pressure regulator, or a Mach-effect drive monitor. Any new loop must:

1. Have a bandwidth ≥ 5× away from every existing loop.
2. Pass the same 8 dB / 35° margin spec at nominal and at all 1σ-tolerance corners.
3. Add no more than 5° of phase loss to any existing loop at its crossover.

If a candidate loop fails any of these, it does not ship until the plant model is updated and the bandwidth ladder is renegotiated. No exceptions; the controls discipline is what protects the science.

9. Open questions

1. **Nonlinear regime.** All margins above are linear. The cavity bistability (Paper 06) is nonlinear; we operate well below threshold but a worst-case drift could cross it. Captured as a runbook gate G-CTRL-04.
2. **Time-varying plant.** The cold-mass temperature drift over a 12-hour run shifts pole locations by ~3 %. Within margin but visible in the in-flight estimator.
3. **Observability.** The phonon population n_b is observable only through the homodyne. A homodyne fault means we are flying L1, L2 blind on n_b — the safety FSM treats this as a trip condition exactly because of that.

Next paper: 22 — Bayesian Analysis Pipeline & Evidence Thresholds.

Paper 22 — Bayesian Analysis Pipeline and Evidence Thresholds

> Status: methodology paper, v0.5. Reading order: assumes 04-falsification-and-roadmap, 09-metrology, 14-safety-and-blinding. Companion code: src/inference.ltl — sealed priors, Bayes-factor calculator, look-elsewhere correction, evidence ledger.

1. Why this paper exists

Paper 09 set the SNR target. Paper 14 set the blinding rails. This paper says, in unambiguous numerical terms, **what counts as evidence**. Without it, "we saw something" is meaningless: any sufficiently flexible model will fit any sufficiently noisy data. We pre-commit to a Bayesian decision rule with sealed priors, sealed thresholds, and a sealed look-elsewhere correction. Past that point, the answer is what the arithmetic says.

The headline numbers, set here and registered in the v0.5 tag:

> **Discovery threshold:** $\ln B_{10} \geq 11.5$ ($\approx 5\sigma$ equivalent for our nuisance dimensionality), after the look-elsewhere correction of §6. **Hint threshold:** $4.6 \leq \ln B_{10} < 11.5$ ($\approx 3\sigma$). **Null:** $\ln B_{10} < 0.7$ (decisive against signal under our priors, $B_{01} > 2$).

These are not negotiable post-hoc. The blinding key (Paper 14) is released only after the analyst has committed to the threshold band their result lands in.

2. The two competing hypotheses

H_0: shot-noise-limited homodyne output, no signal at the predicted sideband location.

H_1: shot-noise-limited homodyne output **plus** a signal at the predicted sideband location with amplitude a and quadrature phase θ as predicted by the optomechanical model (Paper 06).

Critically, H_1 is **a single-point hypothesis in frequency**: the sideband location is fixed by the on-board clock tree (Paper 08), not fitted from the data. The amplitude a and phase θ are free, but only inside the priors of §3.

The likelihood, given a homodyne time-series $\mathbf{y}$ of length N samples after the DSP pipeline (Paper 16):

$$p(\mathbf{y} | H_1, a, \theta, \boldsymbol{\nu}) = \prod_{k=1}^{N} \mathrm{N}\left(y_k - a\cos(2\pi f_{\mathrm{sb}} t_k + \theta);\ \sigma_k(\boldsymbol{\nu})\right)$$

with $\boldsymbol{\nu}$ the nuisance vector (laser-power drift, temperature drift, residual phase noise; nine parameters, see §4). The shot-noise variance σ_k depends on the instantaneous photodetector current and is estimated per-window from the out-of-band spectrum.

3. The priors — sealed before the run

Sealed in src/inference.ltl under `SEALED_PRIORS_V0_5`, hash committed to the v0.5 tag.

Signal amplitude a (units: V at homodyne output):

$$p(a) = \text{HalfNormal}(\sigma_a = 5.0 \times 10^{-9})$$

The scale σ_a is set so that the median predicted Mk1 amplitude (per Paper 06) is at the prior median. We deliberately **do not** use a log-uniform prior over many orders of magnitude; that would overweight implausibly large signals and inflate B_{10} spuriously.

Signal phase θ:

$$p(\theta) = \text{Uniform}(0, 2\pi)$$

Uniform; we have no reason to prefer any quadrature.

Nuisance vector $\boldsymbol{\nu}$: Gaussian priors centred on the calibration values of the run, with widths from the v0.3 prototype's characterisation (Papers 12, 15, 16). Sealed values in src.

4. The nuisance dimensions

The nine nuisance parameters that enter $\boldsymbol{\nu}$:

#	Parameter	Symbol	Prior (Gaussian)
1	Laser power drift	$\Delta P_L/P_L$	0 ± 10^{-3}
2	Cavity-frequency thermal drift	Δf_c	0 ± 200 kHz
3	Mechanical-frequency drift	$\Delta\Omega_m$	0 ± 2 kHz
4	Homodyne gain drift	$\Delta g/g$	$0 \pm 5 \times 10^{-4}$
5	LO phase drift	$\Delta\phi_{\text{LO}}$	0 ± 10 mrad
6	Cold-mass T error	ΔT	0 ± 5 mK
7	Vibration platform residual	σ_{vib}	half-Gaussian, 0.5 pm
8	Detector dark-current	i_{dark}	half-Gaussian, 50 pA
9	DC offset of homodyne output	V_{DC}	$0 \pm 1\,\mu$V

These are marginalised numerically (importance sampling, 10^5 draws per evidence calculation) — the Bayes factor itself, not just a maximum-likelihood number.

5. The Bayes factor — operational definition

$$B_{10} = \frac{p(\mathbf{y}|H_1)}{p(\mathbf{y}|H_0)} = \frac{\int da\, d\theta\, d\boldsymbol{\nu}\; p(\mathbf{y}|H_1, a, \theta, \boldsymbol{\nu})\, p(a)\, p(\theta)\, p(\boldsymbol{\nu})}{\int d\boldsymbol{\nu}\; p(\mathbf{y}|H_0, \boldsymbol{\nu})\, p(\boldsymbol{\nu})}$$

Computed by nested sampling (dynesty), 2 000 live points, until the evidence stops moving by more than 0.1 in $\ln Z$. We log $\ln Z_1$, $\ln Z_0$, and the difference $\ln B_{10}$ separately so that an audit can re-derive the result from the same data and the same sealed priors.

6. Look-elsewhere — the explicit correction

We test only **one** sideband location per run, fixed by the on-board clock — so there is genuinely no LEE within a single run. The LEE arises across the run *campaign*: if we run the experiment N_{run} times with independent thermal/mechanical drifts, the chance of a spurious B_{10} excursion grows.

Pre-registered correction:

$$\ln B_{10}^{\text{discovery}} = 11.5 + \ln N_{\text{run,planned}}$$

For the year-1 plan of $N_{\text{run,planned}} = 50$ runs (Paper 04):

$$\ln B_{10}^{\text{discovery}} = 11.5 + \ln 50 \approx 15.4$$

This is the **post-LEE** threshold. The campaign-wide false-discovery rate at this threshold is < 1%.

If a single run lands in the hint band (4.6 – 11.5) the campaign is extended, with the threshold for discovery being raised accordingly per the same formula (Paper 04 §5 specifies the extension protocol and the abandonment criteria).

7. The decision tree

The pipeline is a strict pre-committed sequence. Each step has exactly one input and one output; no step may be skipped, reordered, or re-entered.

1. **Run completes.** DAQ + DSP emit the time-series $y[k]$ and the per-window shot-noise estimate $\sigma[k]$ (Paper 16).
2. **Blinding still on.** The analyst inspects the data only as summary statistics that cannot resolve $\ln B_{10}$ to better than ±2. The analyst commits — in writing, hash-chained — to which of the four bands they expect the result to land in *before* the blind is touched.
3. **Compute $\ln B_{10}$.** Sealed priors loaded; nested sampling runs to convergence; the result is appended to the evidence ledger.
4. **Classify against the four bands** (with $L^* = 11.5 + \ln N_{\text{run,planned}}$):

   - $\ln B_{10} < 0.7$ — **Null.** Logged, blinding kept on, the run is published as a null result in the campaign roll-up.
   - $0.7 \le \ln B_{10} < 4.6$ — **Inconclusive.** Logged, blinding kept on, the run is rolled into the next batch and the campaign is extended by one run.
   - $4.6 \le \ln B_{10} < L^*$ — **Hint.** Blinding key released for this run only; the result is logged and the campaign is extended per Paper 04 §5.
   - $\ln B_{10} \ge L^*$ — **Discovery.** Blinding key released; the replication contract of §8 is invoked.

The classifier function is `classify()` in src/inference.ltl; it takes a single $\ln B_{10}$ and returns the band, with the threshold table sealed at the v0.5 tag and verifiable by hash.

8. The replication contract

A single discovery threshold crossing is **not** a discovery. The project commits, in advance, to:

1. The blind is opened only on the first crossing.
2. Two further independent runs (different cavity cell, different operator, different week) must each return $\ln B_{10} \geq 4.6$ under the same sealed priors and re-sealed nuisance posteriors.
3. The data, the sealed priors, the analysis code, and the nested-sampling chains for all three runs are released simultaneously to a public archive.
4. Independent replication by another group is requested before any claim of discovery is made externally. The internal threshold crossing alone authorises the data release, not the claim.

If step 2 fails, the original crossing is logged as a probable look-elsewhere artefact and the campaign continues with the threshold for discovery raised by $\ln 3$ (the multiple-test penalty for the three-run replication attempt itself).

9. What this paper forbids

The whole point of pre-registration is to make later flexibility illegal. After the v0.5 tag the following are not permitted:

- Changing any prior in §3 or §4. New priors require a new tag and a fresh campaign.
- Changing the threshold $L^* = 11.5 + \ln N_{\text{run,planned}}$.
- Changing the test statistic from $\ln B_{10}$ to anything else.
- Selecting which runs enter the campaign analysis on the basis of their data (“the cryostat was a bit warm that day”).
- Reporting any sub-component of $\ln B_{10}$ (e.g. the signal-amplitude posterior alone) as a result.

The hash chain (Paper 14) records every analysis call. An attempt to bypass these forbids the hash to verify, and the integrity check that ships with every public release will flag it.

10. Open questions

1. **Calibration of nuisance priors.** Several priors in §4 come from v0.3 prototype data; Mk1 first-light may show that a nuisance distribution is not Gaussian. We commit, *before* any science, to running 10 calibration runs at zero excitation and re-sealing the nuisance priors based on those before any science runs are analysed. This is a one-time recalibration.
2. **Model-space extension.** Should our H_1 family be enlarged to include alternative signal shapes (e.g. broadband)? Currently no — the model-selection prior cost would

dominate. If a hint ever surfaces, we extend the model only as a separate v0.x tag with its own sealed priors.

3. **External calibration.** We have not yet sourced an independent reference of comparable systems' Bayes-factor pipelines. The GW community's `pyCBC` and `bilby` provide good templates; a v0.6 cross-check is on the roadmap.

See 00-index for the corpus map.

Paper 23 — Lagrangian Formulation: Optomechanical + Mach-Effect Coupling

> Status: theoretical paper, v0.6. Reading order: assumes 01-foundations, 06-cavity-optomechanics, and 07-mach-effect. Companion code: src/lagrangian.ltl — symbolic plant model, exports linearised state-space matrices used by src/control.ltl.

1. Why a Lagrangian, and why now

Papers 06 and 07 derived the cavity and the speculative propulsion term separately. They live in different chapters of the textbook: one in quantum optics, the other in Sciama-style relational mechanics. To make a falsifiable prediction we need a *single* action that yields both the homodyne sideband signal and the candidate force, by separate Euler–Lagrange equations of the same fields. This paper writes that action down, derives the noise-equivalent force the Mk1 testbed can resolve, and shows where every term in the Mk1 firmware plant model comes from.

Without this paper the project is two unrelated theories sharing a cryostat; with it, the prediction is one number derivable from one action.

2. Fields and degrees of freedom

The full system has six fields/coordinates:

Symbol	Object	DoF
$\hat{a}, \hat{a}^\dagger$	optical cavity mode at ω_c	1 (complex)
$\hat{b}, \hat{b}^\dagger$	mechanical mode of the photonic crystal at Ω_m	1 (complex)
$\hat{c}, \hat{c}^\dagger$	RF drive carrier at $\omega_d \approx \Omega_m$	1 (complex)
$X_\mu(\tau)$	hull centre-of-mass world-line	4
$\phi(x)$	Sciama scalar (long-range gravitational potential)	1
$\Phi_{\text{em}}(x)$	electromagnetic four-potential	4

The first three are quantised; the last three are treated as classical backgrounds at Mk1 sensitivity. We will quantise ϕ in a separate paper if Mk1 first-light produces a hint above $\ln B_{10} \geq 4.6$ (Paper 22).

3. The action

We split the action into four pieces:

$$S = S_{\text{opt}} + S_{\text{mech}} + S_{\text{int}} + S_{\text{Mach}}$$

3.1 Optical part

A driven, damped cavity mode in a rotating frame at the drive frequency ω_d (so the bare term is the detuning $\Delta = \omega_c - \omega_d$):

$$S_{\text{opt}} = \int dt\ [i\hbar\, \hat{a}^\dagger \dot{\hat{a}} - \hbar\Delta\, \hat{a}^\dagger \hat{a} + i\hbar\sqrt{\kappa_{\text{ext}}}\,(\hat{a}^\dagger \alpha_{\text{in}} - \hat{a}\alpha_{\text{in}}^*)]$$

with $\kappa = \kappa_{\text{ext}} + \kappa_0$ the total cavity linewidth (paper 05) and α_{in} the coherent input amplitude.

3.2 Mechanical part

A high-Q phonon mode of the zipper cell:

$$S_{\text{mech}} = \int dt\, \hbar\Omega_m\, (\hat{b}^\dagger \hat{b} + \tfrac{1}{2}) + \frac{m_{\text{eff}}}{2}\, \dot{X}_{\text{cm}}^2$$

The second term is the rigid-body kinetic energy of the hull centre-of-mass; the optomechanical mode and the rigid-body motion share no kinetic mixing at this order because the zipper mass is $10^6\times$ smaller than the hull mass.

3.3 The standard optomechanical interaction

This is the textbook radiation-pressure coupling (paper 06):

$$S_{\text{int}} = -\int dt\, \hbar g_0\, \hat{a}^\dagger \hat{a}\, (\hat{b} + \hat{b}^\dagger)$$

with $g_0/2\pi = 220$ kHz the per-photon coupling. Linearising about the cavity coherent amplitude $\alpha = \langle \hat{a} \rangle$ gives the linearised coupling rate $G = g_0|\alpha|$, which sets the sideband ladder of paper 06.

3.4 The Mach-effect term

The single explicit speculative term (paper 07). Following Woodward, the effective inertia of an accelerating, energy-storing object depends on the gravitational potential of the rest of the universe. For a stored optomechanical energy density $u(x,t) = \hbar\Omega_m n_b(x,t)/V$ — the cavity's instantaneous phonon energy per unit volume — the additional Lagrangian density is:

$$\mathsf{L}_{\text{Mach}} = \frac{1}{4\pi G \rho_0 c^2}\, \partial_t u\, \partial_t \phi$$

with ρ_0 the mean cosmological mass density and ϕ the Sciama scalar at the cavity location. Integrating by parts and using $\nabla^2\phi \approx 4\pi G\rho_0$ in the local universe (Sciama's identification), the variation with respect to X_{cm} yields a force proportional to $\partial_t(\partial_t u)$:

$$F_{\text{Mach}} = \frac{1}{c^2 \rho_0}\, \partial_t^2 \langle u \rangle \cdot \rho_{\text{hull}} V_{\text{hull}}$$

The sign and pre-factor depend on conventions (Sciama 1953, Woodward 2013, Tajmar 2021); the magnitude does not. For the Mk1 parameters of paper 06 this gives a peak force of order 10^{-9} N at $\Omega_m/2\pi = 1.5$ GHz drive.

4. Equations of motion — the four physical predictions

Varying the action with respect to each field gives four coupled equations. We list them in operator form; the c-number versions used by the firmware solver are in src/lagrangian.ltl.

Optical mode (Heisenberg–Langevin, with input fluctuations $\hat{a}_{\text{in}}$):

$$\dot{\hat{a}} = -\left(\frac{\kappa}{2} + i\Delta\right)\hat{a} + ig_0\hat{a}(\hat{b} + \hat{b}^\dagger) + \sqrt{\kappa_{\text{ext}}}\,\hat{a}_{\text{in}}$$

Mechanical mode (with thermal bath at temperature T_{bath}):

$$\dot{\hat{b}} = -\left(\frac{\gamma_m}{2} + i\Omega_m\right)\hat{b} + ig_0\hat{a}^\dagger\hat{a} + \sqrt{\gamma_m}\,\hat{b}_{\text{in}}$$

Hull centre-of-mass (classical, the Mach term is the speculative addition):

$$m_{\text{hull}}\ddot{X}_{\text{cm}} = F_{\text{Mach}}(t) + F_{\text{ext}}(t)$$

Sciama scalar (Poisson, sourced by stored energy density):

$$\nabla^2\phi = 4\pi G\rho_0 + \frac{4\pi G}{c^2}\,u(x,t)$$

5. The integrated prediction — one number

Solving the coupled system in steady state under sinusoidal drive at $\omega_d = \Omega_m$, the predicted Mk1 thrust amplitude is:

$$F_{\text{Mach,pred}} = \frac{2\hbar G\Omega_m^2|\alpha|^2 g_0^2}{c^4\rho_0\kappa^2}\cdot\chi_g(\Omega_m)$$

where $\chi_g(\Omega_m)$ is a dimensionless geometry factor of order unity (it captures the cavity mode-volume / hull-volume overlap). For Mk1: $|\alpha|^2 = 10^9$ photons, $g_0/2\pi = 220$ kHz, $\kappa/2\pi = 1.5$ GHz, $\Omega_m/2\pi = 1.5$ GHz, $\rho_0 = 9.7\times10^{-27}$ kg/m³ (cosmological mean):

$$F_{\text{Mach,pred}} \approx 2.4\times10^{-12}\ \text{N}$$

This is the **central prediction** of the project. It must show up both as an in-band homodyne sideband (paper 06) and as a load-cell reading on the hull (paper 03 §6). Either signal alone is necessary but not sufficient; both together cross the discovery threshold of paper 22 only if the *amplitude ratio* between the optical and mechanical channels matches §3.4 ± 30 %.

6. The noise-equivalent-force derivation

A force F on the hull translates to a phase shift on the homodyne:

$$\delta\phi_{\text{hom}}(\omega) = \frac{2G\,F(\omega)}{m_{\text{hull}}\,\omega^2\,\hbar\,\kappa\,\sqrt{n_b + 1/2}}$$

The shot-noise-limited phase floor is $\delta\phi_{\text{shot}} = 1/\sqrt{2|\alpha|^2}$ (paper 15). Solving for the equivalent force:

$$F_{\text{NE}} = \frac{m_{\text{hull}}\,\omega^2\,\hbar\,\kappa}{2G\,|\alpha|\,\sqrt{2(n_b + 1/2)}} \cdot \frac{1}{\sqrt{T_{\text{int}}}}$$

For Mk1 with 1-hour integration: $T_{\text{int}} = 3600$ s, $|\alpha| = 3 \times 10^4$, $n_b \approx 1.0$ (paper 10 ground-state cooling), $m_{\text{hull}} = 125$ kg, $\omega = \Omega_m$:

$$F_{\text{NE}} \approx 8 \times 10^{-15}\ \text{N}/\sqrt{\text{Hz}} \;\Rightarrow\; F_{\text{NE,1h}} \approx 1.3 \times 10^{-16}\ \text{N}$$

We resolve the predicted 2.4×10^{-12} N at SNR $\approx 1.8 \times 10^4$ in one hour. **The signal is not noise-limited; it is systematic-limited.** This is exactly why Papers 14, 19, 20, 22 exist — they are the systematic-error contract.

7. Where each plant-matrix coefficient comes from

The firmware (src/control.ltl) uses a six-state linearised plant. Every entry of A, B, G is traceable to a term in §3:

A_{ij}	**Linearisation of**	**Comes from**
$\partial\dot{\phi}/\partial\delta f_L$	$\dot{a}$	S_{opt} rotating-frame term
$\partial\dot{n}_b/\partial\delta f_R$	$\dot{b}$	S_{int}, linearised
$\partial\delta T/\partial n_b$	hull energy balance	$\hbar\Omega_m$ deposited per phonon
$\partial\delta f_L/\partial\delta T$	thermo-optic	cavity refractive index $\partial n/\partial T$
$\partial\delta f_L/\partial\delta z$	strain coupling	zipper geometry, paper 17
$\partial F_{\text{Mach}}/\partial n_b$	L_{Mach}	second time-derivative of u

That last row is the *only* row of the plant that depends on the speculative term. If the Mach-effect term is wrong, that row's coefficient is zero and we measure no force. If it is right, the coefficient takes the value derived in §5. **Mk1 measures that coefficient directly** by sweeping the drive amplitude $|\alpha|^2$ and looking for the predicted quadratic scaling with stored energy.

8. What this paper does *not* claim

1. We do **not** claim the Mach-effect term in §3.4 is the only or correct extension of GR. Other candidates (Hoyle–Narlikar, Mashhoon–Theiss, Wagoner–Will) give terms with

different power spectra. The paper 22 prior is wide enough to encompass several; a positive result narrows the model space, it does not prove §3.4.

2. We do **not** claim ρ_0 is the right scaling parameter. If the relevant background is local rather than cosmological, the prediction in §5 changes by orders of magnitude. We log the measured F_{Mach} and the measured $|\alpha|^2$ scaling; the *ratio* of the two is the parameter-free prediction.
3. We do **not** claim Lorentz invariance of the Mach term. The commissioning runbook (paper 18) deliberately runs at four sidereal times to test for a sidereal modulation; absence of one constrains the term's preferred-frame structure.

9. Open questions, in order of importance

1. **Quantisation of ϕ.** At Mk1 sensitivity the Sciama scalar is classical. At Mk3 (paper 03), zero-point fluctuations of ϕ may set the noise floor and need quantising. Out of scope here; captured as a v1.0 paper.
2. **Back-reaction on the cavity.** A force on the hull moves the cavity in the lab frame, which in turn shifts ω_c via the tether strain (paper 17). The plant model includes this loop at first order; a second-order correction is in the open-questions list of paper 21.
3. **Coupling sign.** §3.4 takes a specific sign convention. A wrong sign predicts a force *against* the radiation pressure rather than with it; the homodyne signal would be 180° out of phase with the load-cell signal. Mk1 measures the relative phase to ±2°, so this is a clean test even if the magnitude prediction is off.

Next paper: 24 — Reproducible Builds and Signed-Firmware Audit Trail.

Paper 24 — Reproducible Builds and the Signed-Firmware Audit Trail

> Status: methodology paper, v0.6. Reading order: assumes 11-fabrication-qc, 14-safety-and-blinding, and 22-bayesian-pipeline. Companion code: src/provenance.ltl — attestation chain, signed-artifact verification, ledger anchor.

1. Why this paper exists

Every prior anomalous-propulsion claim that drew serious scrutiny shared a forensic problem: **the experimenters could not show, after the fact, exactly what software was running when the result was recorded.** Firmware was reflashed between runs, build environments drifted, plotting scripts were not version-controlled, and the post-publication conversation degenerated into "but in our build…". This paper specifies the contract that makes that conversation impossible for element-115-drive: every byte of firmware, every byte of analysis code, and every byte of raw data is reproducible from public sources, signed, and tied to the science result by hash.

The claim of this paper:

> A reviewer with the public sources, the public sealed priors (paper 22), and the SHA-256 of a published result can — on commodity hardware — re-derive the firmware binaries that were running, the analysis pipeline that ran, and within numerical tolerance the analysis output, byte-for-byte for the firmware and bit-for-bit for the result hash.

If that claim ever fails to hold, the result is withdrawn until it does.

2. The reproducibility contract — five layers

#	Layer	Tool	Verifier acceptance
1	Source archive	git, signed tags	tag signature verifies against the project key
2	Build environment	nix-pinned toolchain	flake.lock SHA matches; toolchain hash matches
3	Firmware binary	reproducible build	byte-identical reproduction of every `.elf` and `.bin`
4	Run-time attestation	TPM-backed signing	every ledger entry signed by an in-firmware key whose hash is in §3
5	Analysis output	sealed pipeline	re-run on the published raw data yields the same $\ln B_{10}$ ± numerical tolerance

A reviewer who cannot verify any one layer rejects the result. A reviewer who can verify all five accepts it as *internally consistent*; whether it is also *physically true* is a separate question that paper 22 addresses.

3. Layer 1 — source archive

The git repository is the source of truth.

- The default branch is fast-forward only. No history rewrites, no force pushes, no rebase merges.
- Every release is a git tag, signed with the project's Ed25519 key. The public key is committed at `keys/release.pub` and pinned in this paper as:

``` ed25519:AAAAC3NzaC1lZDI1NTE5AAAAIB6e1nT4AB

---

fingerprint: SHA256:Wd2pHo3KbV2H4xLs+8ePBT1n0Z7QqKzN5fA8aWmYqCk ```

Reviewers should verify the fingerprint against this paper's signed PDF (paper 24, v0.6 tag).

- Every release ships with an SBOM (`.sbom.spdx.json`) generated from the nix flake, pinning every transitive dependency by hash.
- Release notes are themselves committed under `release-notes/` and signed; no out-of-band release announcement exists.

## 4. Layer 2 — build environment

We use Nix (https://nixos.org) flakes for the build environment. Specifically:

- The toolchain is pinned in `flake.lock` to specific git revisions of nixpkgs. There are no `*` or `latest` references anywhere in the build configuration.
- The Lateralus compiler (`ltlc`), the assembler, the linker, and every Python tool used in CI are *each* pinned by a content hash, not a version number.
- Build hosts are stateless: every CI run starts from a freshly-derived nix store. No host filesystem state leaks into the build.
- The build is timestamp-clean: `SOURCE_DATE_EPOCH` is set from the git commit time of the tag; all build outputs use that timestamp.

A reviewer can reproduce the toolchain on their own machine:

```
git clone https://lateralus.dev/git/element-115-drive.git
cd element-115-drive
git verify-tag v0.6
nix develop   # drops into the pinned toolchain
make firmware-all   # builds every src/*.ltl module
```

The expected output of `make firmware-all` is recorded in `build/expected-hashes.json`. If any artefact's SHA-256 differs from the recorded value, the reviewer has reproduced something else — and we want to know what.
```

5. Layer 3 — firmware binary

A reproducible build is one where, given the same source and the same toolchain, the output is byte-identical regardless of who ran the build, on what host, at what time. We enforce this by:

1. Compiler determinism. The Lateralus compiler is single-threaded in `--reproducible` mode and emits no debug paths or absolute filenames into the binary.
2. Linker determinism. Symbol order is sorted by symbol name, not by link order. Padding bytes are zero, not undefined.
3. Build path stripping. Source paths in DWARF debug info are rewritten to `/build/<module>/<file>` regardless of host path.
4. Embedded assets ship pre-compiled. There is no on-build timestamping; the build date is the tag date.

Every `.elf` produced by `make firmware-all` has its SHA-256 recorded in `build/expected-hashes.json`. The CI pipeline rebuilds on three independent hosts (Linux x86_64, Linux aarch64, NetBSD x86_64) and asserts the hashes match. If they don't, the build is broken and the release is blocked.

6. Layer 4 — run-time attestation

The science µC has an STMicro STSAFE-A110 secure element. At first boot it generates an Ed25519 keypair internal to the chip; the public key is read out, signed by the project release key, and written to a fuse-locked region of the µC. From that point on:

- Every entry in the safety FSM hash-chain (paper 14) is signed by the µC's secure-element key.
- Every entry in the energy ledger (paper 19), the EMC scan ledger (paper 20), the margin ledger (paper 21), and the evidence ledger (paper 22) is signed by the same key.
- The project release key signs the per-device public key; a chain of trust from a science result back to a specific physical µC is therefore verifiable on a reviewer's laptop using only public data.

The chain that must verify, in order:

1. `release_pub` (in this paper) signs `device_pub` (in `devices/<serial>.pub`).
2. `device_pub` signs every ledger entry.
3. The ledger's hash-chain (paper 14 §4) links each entry to the previous one; tampering with any single entry breaks the chain from that point forward.
4. The final ledger entry's hash is published as part of the science result; it is what readers cite.

Anyone with the public release key, the device pubkey, and the ledger can verify the entire run end-to-end. They do *not* need our cooperation; they do *not* need our infrastructure; they only need this paper, the published code, and the published ledger.

7. Layer 5 — analysis output

The analysis pipeline (paper 22) is a single sealed module src/inference.ltl. Reproducing a published $\ln B_{10}$ requires:

1. The published raw time-series and shot-noise estimate (in `runs/<run_id>/y.npy`, `runs/<run_id>/sigma.npy`).
2. The sealed priors hash from this paper's tag (paper 22 §SEAL_HASH).
3. The nested-sampling parameters: `n_live = 2000`, `tol_lnz = 0.1`, plus the random seed published in the run report.

Running:

```
make verify-run RUN=<run_id>
```

invokes the same `inference::compute_evidence` that the live pipeline ran. Output is compared to the published $\ln B_{10}$ value; they must agree to within the nested-sampling tolerance.

Numerical tolerance is non-zero because nested sampling is stochastic. The tolerance is:

$$\left| \ln B_{10}^{\text{live}} - \ln B_{10}^{\text{verify}} \right| \le 0.2$$

which is the 95-th percentile of the seed-to-seed spread we measured on synthetic data with 2 000 live points. A discrepancy outside this band is a defect of either the verifier's pipeline or our published data, and is treated as a withdrawal-grade incident.

8. The published-result manifest

Every run report ships with a manifest JSON containing:

- `run_id`, `iso_8601_started`, `iso_8601_ended`
- `device_pub` and `device_pub_signature`
- `firmware_hashes` — every running `.elf`'s SHA-256
- `ledger_head` — final hash of the run's ledger
- `evidence_report` — `ln_z0`, `ln_z1`, `ln_b10`, verdict, seed
- `corpus_tag` — the git tag of the paper corpus the run was conducted under (e.g. `v0.6`)
- `manifest_signature` — Ed25519 signature of all the above

A reviewer rebuilds the manifest from public sources and checks every signature. If any check fails, the run is rejected.

9. Counter-attack surface

We considered, and have a mitigation for, each of the following:

Attack	Mitigation
“Reflash the firmware between runs”	Layer 4 attestation: every ledger entry is signed by the on-chip key whose pub is in the manifest
“Cherry-pick which runs to publish”	Layer 4 ledger is append-only; missing run IDs are detectable; runs are numbered consecutively from cryostat first-cool
“Tamper with the analysis output”	Layer 5 reproducibility: anyone re-runs and checks
“Tamper with the raw data”	Raw data is committed to the ledger by hash before analysis (paper 16); ledger entry is signed
“Substitute a different µC”	Project release key signs the µC pubkey at provisioning; the signature is in the manifest
“Compromise the project release key”	Key rotation paper (out of scope here); old signatures remain valid; new signatures use a new key tagged in a fresh paper

There is one attack we cannot defend against: the operator lying about what was physically connected to the µC. We mitigate it *socially* — by inviting external observers (paper 04 §6) — and *physically*, by photographing the rack at gates G-PRE-01 and G-RUN-99 of the runbook (paper 18). Beyond that we accept a residual trust-the-operator risk; that risk is what external replication addresses.

10. Why this is unusually paranoid

A reviewer reading this paper might object that no other physics experiment imposes this much paranoia on its own software stack. The reasons we do are specific to *this* project:

1. The result, if positive, is extraordinary. Extraordinary evidence is the standard. Software-trust is a pre-requisite for any evidence claim.
2. The signal is small enough that a software bug in the analysis pipeline could fake it — see paper 25 for a list of prior claims where exactly that happened.
3. The result, if negative, must close the door cleanly. A forensically tight null is more useful to the field than a forensically loose hint.

The cost of this layer is real: about three engineer-months of work and one part of additional CI infrastructure. It is paid once and pays back forever.

Next paper: 25 — Comparison vs Prior Anomalous-Propulsion Claims.

Paper 25 — Comparison vs Prior Anomalous-Propulsion Claims

> Status: methodology paper, v0.6. Reading order: assumes 04-falsification-and-roadmap, 09-metrology, 22-bayesian-pipeline, 24-reproducible-build. No companion code (methodology paper).

1. Why we have to write this

Every claim of anomalous propulsion in the last fifty years has collapsed under one of three failure modes: irreproducible build chain, unblinded analysis, or a thermal/EMI artefact mistaken for signal. This paper goes through the most prominent claims, explains which failure mode killed each, and shows the structural rules in this project that make the same failure mode impossible by construction. The point is not to disrespect prior researchers; several of them did real, careful work that was undone by infrastructure problems. The point is to inherit their lessons explicitly so we don't repeat them.

This paper is the one a sceptical funder, journalist, or peer reviewer should read first. If they remain unconvinced after reading it, none of the technical papers will help. If they are convinced, the rest of the corpus is a roadmap to the test that either confirms or kills the claim within twelve months.

2. The five claims we anchor against

Claim	Year	Headline result	Status by 2025
Lazar / element-115 reactor	1989	"Gravity wave amplifier", anti-matter reactor	No experimental data ever released; no peer-reviewed paper
Podkletnov / superconducting disc	1992	0.05 % weight reduction over rotating YBCO disc	One unblinded replication; multiple blinded null replications
Shawyer / EmDrive	2003	μN thrust from RF cavity, no propellant	NASA-Eagleworks 2016 hint; multiple blinded null replications by 2018
Tajmar / artificial gravitational fields	2006–2021	gravitomagnetic effect from rotating cryogenic ring	Tightened upper limits each campaign; never crossed 5σ
Woodward / MEGA drive	2014–	Mach-effect thrust from PZT stack	Replication ongoing; Tajmar 2021 set tight upper limits

3. Per-claim post-mortem

3.1 Lazar (1989)

What was claimed: a gravity-wave amplifier powered by element-115, generating an "out-of-phase" gravitational distortion field for propulsion.

What was actually shown: nothing measurable, by anyone, ever. The account is interview-only, the device was reportedly classified, and no apparatus, data, or independent measurement has been published in the intervening 36 years.

Failure mode: **complete absence of falsifiability**. There is no test that could disprove the claim because there is no specified apparatus.

How element-115-drive avoids this:

Lazar problem	Our structural answer
No specified device	Papers 03, 05, 06, 17 specify every dimension to ±tolerance
No specified test	Paper 04 specifies the abandonment criteria
No data release	Paper 24 specifies the manifest format and signing chain
No independent measurement	Paper 04 §6 mandates external observation at three gates

Our project name retains the “115” digits as historical homage. We do not claim element 115 is involved in any of the physics; we *explicitly drop* that framing in paper 01. We are not building Lazar’s claim; we are building a falsifiable test of a different claim using inheritance only of the cavity-driven framework.

3.2 Podkletnov (1992)

What was claimed: a 0.05 % weight reduction for objects suspended above a rotating, RF-pumped, superconducting YBCO disc at 5000 rpm.

What was actually shown: a single laboratory’s measurement, with the test mass on a balance directly above the rotor, at room temperature. The result was unblinded; the analysis was unsealed; the apparatus produced significant magnetic and thermal fields.

Failure mode: **unblinded analysis with multiple known confounds**. The most cited replication attempts (Tajmar 2002, Woods 2003) found no effect once thermal convection and Lorentz forces on conducting test masses were properly nulled.

How element-115-drive avoids this:

Podkletnov problem	Our structural answer
Unblinded analysis	Paper 14 AES-GCM blinded predictions; paper 22 sealed priors
Thermal convection of test mass	Paper 17 vibration / thermal isolation; paper 10 cryogenic vacuum
Magnetic-force confound	Paper 20 Faraday cage + 3-axis fluxgate at the cavity
Single-lab result	Paper 24 reproducible build + paper 04 external replication contract

The structural difference is paper 22: every threshold and prior is sealed before any data is taken. A Podkletnov-style analysis where the threshold is set after looking at the data is *forbidden by the verifier*, not just discouraged.

3.3 Shawyer / EmDrive (2003) and NASA-Eagleworks (2016)

What was claimed: a closed asymmetric RF cavity producing thrust without expelling reaction mass — explicitly a violation of momentum conservation as currently understood.

What was actually shown: a series of μN-level thrust readings on a torsion balance. The 2016 NASA-Eagleworks paper was published in *JPP*, peer-reviewed.

What killed it: thermal expansion of the cavity walls and the cable harness produced exactly the observed signature; once Tajmar 2018 ran the same apparatus blinded with thermal nulls in place, the signal vanished into shot noise. The cable-harness Joule heating moved the centre of mass of the torsion balance by enough to fake a μN-class force.

Failure mode: **thermal-expansion artefact** confused with a real force, in an unblinded analysis. The cabling pulled the balance, not the cavity.

How element-115-drive avoids this:

EmDrive problem	Our structural answer
Cable-harness thermal pull	Paper 03 §6 fibre-optic-only signal entry to the test mass
Single-channel readout	Paper 06 + paper 23 require correlated optical and load-cell signals; either alone is insufficient
Unblinded analysis	Paper 14 + paper 22 sealing
Drive-power-correlated artefacts	Paper 20 G-EMC-03 explicitly checks for drive-power-correlated detector responses with the cavity disconnected

The deepest structural answer is the **dual-channel requirement**: paper 23 §5 predicts both an optical sideband and a mechanical force from the same Lagrangian. A claim is taken seriously only if both appear with the predicted phase relationship. EmDrive had no such correlated cross-channel; if it had, its thermal artefact would have been detected at first analysis.

3.4 Tajmar (2006–2021)

What was claimed (early): an artificial gravitomagnetic field from rotating, cryogenic superconducting rings, large enough to move nearby accelerometers.

What was actually shown: tightening upper limits each campaign; each generation of Tajmar's apparatus reduced the residual signal, and by ~2010 the signal vanished into instrument noise.

What this is: **a positive example of the discipline**. Tajmar's later campaigns are exactly the falsifiability model we want to emulate. His 2021 review of MEGA-drive results is a model methodology paper. We diverge from him only in two places:

1. We seal the analysis pipeline ahead of time (paper 22). Tajmar sealed the apparatus but not the analysis.
2. We use a cavity-optomechanical readout, which has a $\sim 10^6$ stronger transfer function from a candidate force to a detector signal than a torsion-balance readout.

If our Mk1 returns a null and Tajmar's later campaigns also return null, we and he are saying the same thing in different ways.

3.5 Woodward / MEGA drive (2014–)

What was claimed: a Mach-effect thrust from a piezoelectrically driven, energy-storing PZT stack, predicted by a Sciama-style relational mechanics extension.

What was shown: small thrust signals on torsion-balance and hanging-pendulum apparatus, replicated by a second lab (Hatch), upper-limited by Tajmar 2021 to below the predicted level.

Failure mode: **borderline reproducible** — different apparatus generations gave different signs and amplitudes; the analysis was not sealed; the prediction depended on a constant (ρ_0) whose value was sometimes set by the data.

How element-115-drive differs:

Woodward problem	Our structural answer
Sign and magnitude shifted between apparatus	Paper 23 fixes the sign and magnitude *before* the run
ρ_0 effectively floated	ρ_0 is fixed at the cosmological mean in paper 23 §3.4; deviation triggers a replication step, not a fit
PZT stack: many simultaneous mechanical modes	Optomechanical zipper has one well-characterised mode at Ω_m; paper 06
Torsion-balance only	Cross-channel optical + mechanical (paper 23 §7)

We *retain* Woodward's Lagrangian term verbatim (paper 23 §3.4). What we add is the cavity-optomechanical channel, which lets us cross-check the prediction with an independent observable.

4. The lessons we encoded as structural rules

From the five post-mortems, six rules. Each is a section of the corpus, not just a lab habit:

Rule	Where it lives
The threshold is sealed before the data are taken	Paper 22 §1, src/inference.ltl `SEAL_HASH`
Every signed result has a reproducible build behind it	Paper 24 layers 1–3
Every result has an attested run-time signature chain	Paper 24 layers 4
The analysis is re-runnable on commodity hardware	Paper 24 layer 5
Two independent observables must agree in phase and magnitude	Paper 23 §5, paper 09 §3
External observers attend at gates G-PRE-01, G-RUN-50, G-RUN-99	Paper 04 §6, paper 18

A claim that does not satisfy all six rules cannot be entered into this project's evidence ledger. It is not a matter of policy; the firmware refuses.

5. What this project will count as a positive result

A positive result is a single chain that the public can audit:

1. The sealed priors of paper 22 are the priors at the v0.6 tag — verified by hash.
2. The firmware running was the v0.6 firmware — verified by the manifest of paper 24.
3. The optical channel returned a sideband at f_{sb} with amplitude a_{opt}, phase θ_{opt}.
4. The mechanical channel returned a load-cell signal at the same sideband location with amplitude a_{mech}, phase θ_{mech}.
5. a_{opt} and a_{mech} agree with the paper-23 §5 prediction to within ±30 %, and θ_{opt} and θ_{mech} agree to ±10°.
6. $\ln B_{10} \geq 11.5 + \ln N_{run,planned}$ (paper 22).
7. Two independent replication runs on different cells, by different operators, at different times, return $\ln B_{10} \geq 4.6$ each (paper 22 §8).
8. The sidereal-time check (paper 23 §8 Q3) returns either a modulation we predict or a null we predict.

Eight conditions, all auditable from public data. Any one missing, the result is logged as a hint and the campaign continues.

6. What this project will count as a negative result

A clean negative is just as publishable, and arguably more useful:

- $\ln B_{10} < 0.7$ on ≥ 40 of the 50 planned runs.
- No drift in either channel correlated with drive amplitude beyond the §6 noise-equivalent floor of paper 23.
- The corpus is published, the data are published, the firmware is published. The conclusion is: *under the conditions specified in this corpus, the Woodward-Sciama Mach-effect Lagrangian term is bounded above by* $X \times 10^{-Y}$ *N at Mk1 sensitivity.*

That is a citable upper limit. It tightens the existing Tajmar 2021 bound by either $\sim 10^3$ in ρ_0-space or $\sim 10^2$ in amplitude space, depending on the analysis cut. Either way it is a real datum.

7. The honest summary

Element-115-drive is **not** a propulsion claim. It is a falsifiable test of one specific Lagrangian term, set up to fail cleanly if the term is wrong and to verify cleanly if it is right. The project's contribution is not "we will fly to Alpha Centauri"; it is "we will either tighten an existing upper limit by three orders of magnitude or we will identify a candidate signal, with a reviewer-verifiable audit trail in either case."

Researchers in the cavity-optomechanics community, the gravitomagnetic community, and the experimental-relativity community are invited to read the corpus, attempt to break the structural argument, and bring their critiques to the public ledger. The project's resilience to those critiques is the project's value, regardless of the eventual sign of the science result.

Next paper: 26 — First-Light Science Campaign.

Paper 26 — First-Light Science Campaign: Pre-Registered Run Plan

> Status: methodology paper, v0.6. Reading order: assumes 04-falsification-and-roadmap, 18-commissioning-runbook, 22-bayesian-pipeline, 24-reproducible-build. Companion code: src/campaign.ltl — sealed per-run hypothesis ledger, run-counter monotonicity guard.

1. The pre-registered campaign

Paper 22 set the analysis discipline. This paper is the **schedule** that discipline runs against. Every run in the year-1 campaign is listed below, with its sealed hypothesis, its predicted $\ln B_{10}$ distribution under each model, and its abandonment trigger. The sealing is enforced by src/campaign.ltl: once the v0.6 tag is committed, no run-list field can be edited without a new paper, a new tag, and a new seal hash.

The campaign is 50 runs across 12 calendar months. It is split into five phases of ten runs each. Phases unlock conditionally on the preceding phase passing its own gate.

Phase	Runs	Purpose	Phase gate
A	1–10	Calibration & nuisance re-characterisation	All 10 nuisance posteriors agree with prototype priors within 1.5σ
B	11–20	Drive-amplitude scaling	$
C	21–30	Sideband-frequency scan	Sweep f_{sb} ±50 MHz; signal must track on-chip clock, not lab clock
D	31–40	Sidereal-time spread	Four runs each at 0h, 6h, 12h, 18h LST; Lorentz-invariance check (paper 23 §8 Q3)
E	41–50	Replication of the strongest band	Whatever band has the highest $\ln B_{10}$ in A–D, replicated by an independent cell

If phase A fails its gate, the campaign halts and a new tag is cut to recharacterise the priors before any other phase begins. This is the single re-sealing the project allows; everything after phase A is sealed at v0.6.

2. Per-run hypothesis sealing

Every run carries five sealed fields, written into the ledger before the cryostat is powered up:

Field	Source	Sealed value (example, run 11)
`run_id`	monotone counter	11
`phase`	from §1	B
`f_sb_target_hz`	paper 08 clock tree	1 500 000 750
`alpha_squared_target`	paper 23 §5 + scan plan	5.0×10^7
`predicted_ln_b10_h1`	nested-sample on synthetic data	9.4 ± 1.1

The sealing function is `campaign::seal_run(run_id);` it appends to the same ledger as the safety FSM (paper 14) and the evidence record (paper 22). The ledger entry signs the tuple of all five fields with the device key (paper 24 §6). After sealing, the run is permitted to start; before sealing, it is not.

A reviewer reading the public ledger sees the predicted $\ln B_{10}$ *before* the run was conducted. They can therefore detect, after the fact, any attempt to retroactively claim a successful prediction.

3. The 50-run schedule (sealed at v0.6)

Phase A — Calibration (runs 1–10)

#	Drive	f_{sb} deviation	Cryostat T	Predicted $\ln B_{10}$ (H_1)	Predicted $\ln B_{10}$ (H_0)
1	OFF	n/a	4.0 K	n/a	-0.1 ± 0.4
2	OFF	n/a	4.0 K	n/a	-0.1 ± 0.4
3	OFF	n/a	4.0 K	n/a	-0.1 ± 0.4
4	OFF	n/a	100 mK	n/a	-0.1 ± 0.4
5	OFF	n/a	100 mK	n/a	-0.1 ± 0.4
6	LOW (10^4)	0	100 mK	1.2 ± 0.5	-0.1 ± 0.4
7	LOW (10^4)	0	100 mK	1.2 ± 0.5	-0.1 ± 0.4
8	LOW (10^4)	0	100 mK	1.2 ± 0.5	-0.1 ± 0.4
9	MID (10^6)	0	100 mK	4.7 ± 0.8	-0.1 ± 0.4
10	MID (10^6)	0	100 mK	4.7 ± 0.8	-0.1 ± 0.4

Phase A gate: nuisance posteriors from runs 1–5 must overlap the priors of paper 22 §4 within 1.5σ on each of the nine dimensions. Failure ⇒ halt, recharacterise, re-seal.

Phase B — Drive-amplitude scaling (runs 11–20)

Predicted $\ln B_{10}$ at H_1 scales roughly with $|\alpha|^4$ (paper 23 §5 squared signal × shot-noise floor). Sweep:

#	$\lvert\alpha\rvert^2$	Predicted $\ln B_{10}$ (H_1)	Predicted (H_0)
11	5×10^7	9.4 ± 1.1	-0.1 ± 0.4
12	5×10^7	9.4 ± 1.1	-0.1 ± 0.4
13	1×10^8	11.7 ± 1.2	-0.1 ± 0.4
14	1×10^8	11.7 ± 1.2	-0.1 ± 0.4
15	3×10^8	14.8 ± 1.4	-0.1 ± 0.4
16	3×10^8		

14.8 ± 1.4 | -0.1 ± 0.4 | | 17 | 1×10^9 | 18.3 ± 1.5 | -0.1 ± 0.4 | | 18 | 1×10^9 | 18.3 ± 1.5 | -0.1 ± 0.4 | | 19 | 3×10^9 | 21.6 ± 1.6 | -0.1 ± 0.4 | | 20 | 3×10^9 | 21.6 ± 1.6 | -0.1 ± 0.4 |

Phase B gate: the *slope* of $\ln B_{10}$ vs $\log |\alpha|^2$, fitted across runs 11–20, must be either consistent with H_0 (slope = 0) or consistent with H_1 (slope ≈ $1.5 \times 2 = 3$ per e-fold, from $|\alpha|^4$). A slope intermediate between these two predictions is *its own* falsification — it indicates an unmodelled drive-correlated artefact and triggers a hold while paper 20's EMC scans are re-run.

Phase C — Sideband-frequency scan (runs 21–30)

Lock the drive at the strongest amplitude that passed phase B. Step f_{sb} across ±50 MHz of the on-chip target. The prediction is that signal tracks the *on-chip clock*, not the lab clock. Specifically, if the lab clock drifts by 5 ppm during a run and the on-chip clock is steered against an internal reference, the signal must follow the on-chip reference. This is the discrimination against a drive-leak artefact (which would track the lab clock).

Per-run predicted $\ln B_{10}$ at the on-resonance run only:

#	Δf_{sb} (MHz)	Predicted (H_1)	Predicted (H_0)
21	0	11.7 ± 1.2	-0.1 ± 0.4
22	+1	9.0 ± 1.0	-0.1 ± 0.4
23	+5	3.1 ± 0.7	-0.1 ± 0.4
24	+20	0.5 ± 0.5	-0.1 ± 0.4
25	+50	0.0 ± 0.4	-0.1 ± 0.4
26	−1	9.0 ± 1.0	-0.1 ± 0.4
27	−5	3.1 ± 0.7	-0.1 ± 0.4
28	−20	0.5 ± 0.5	-0.1 ± 0.4
29	−50	0.0 ± 0.4	-0.1 ± 0.4
30	0	11.7 ± 1.2	-0.1 ± 0.4

Phase C gate: the line-shape, fitted to the ten $\ln B_{10}$ values, must match a Lorentzian centred on the on-chip target with FWHM within ±20 % of the predicted γ_m. A line-shape mismatch triggers a phase 0 mechanical-Q recharacterisation.

Phase D — Sidereal spread (runs 31–40)

Four sidereal-time bins, two runs per bin, plus two repeats of the strongest result so far. The prediction under the Sciama-respecting form of the §3.4 Lagrangian is **no sidereal modulation** above 0.5 % of the on-resonance amplitude. A modulation above that level indicates a preferred-frame coupling and *changes the model*, i.e. the project re-tags as v0.7 and re-seals priors.

#	LST (h)	Repeat-of-strongest?
31, 32	0	no
33, 34	6	no
35, 36	12	no
37, 38	18	no
39, 40	(variable)	yes

Phase D gate: an F-test on the four-bin amplitude variances against the within-bin variance, with a p-value < 0.01 triggering a preferred-frame tag.

Phase E — Independent-cell replication (runs 41–50)

Swap to a second cavity cell from a different fabrication run (paper 11 G7 acceptance). Repeat the strongest five conditions of phases A–D (whichever five gave the highest $\ln B_{10}$). The operator and the analyst rotate so neither is the same as the original run.

Replication thresholds (paper 22 §8): each phase E run must return $\ln B_{10} \geq 4.6$ to confirm the original; failure of two of the five logs the original as a probable artefact.

4. Abandonment criteria

Cumulative across the campaign. If any of the following becomes true *at any point*, the campaign halts and the project enters wind-down per paper 04 §7:

1. ≥ 3 EMC scan failures (paper 20) within any 10 consecutive runs.
2. Any margin-low fault (paper 21 §5) lasting > 1 hour.
3. Cryostat warmup beyond 100 mK during a science run (paper 10).
4. Phase A nuisance check fails after two recharacterisations.
5. Phase B slope falls into the “intermediate” band on two independent fits.
6. Phase C line-shape FWHM is > 3× predicted.
7. Phase E replication: 0 of 5 above $\ln B_{10} = 4.6$.

Each of (1)–(7) is detected by src/campaign.ltl and triggers a sealed `Halt` event in the ledger.

5. The success criterion

Beyond a single discovery threshold crossing (paper 22 §8), the campaign as a whole succeeds if and only if:

- Phase A nuisance posteriors agree with prototype priors.
- Phase B slope is consistent with one of the two pre-registered predictions (H_0 or H_1).
- Phase C line-shape is centred on the on-chip clock.
- Phase D shows the predicted absence of sidereal modulation *or* shows a sidereal modulation with a clean re-tag to v0.7.
- Phase E confirms whatever band phases A–D produced.

If the science verdict is H_0 across the board, that is a **successful campaign with a null result**. The corpus is published, the data are published, and the upper limit on the §3.4 Lagrangian term is a citable result that tightens the existing Tajmar 2021 bound by the factor named in paper 25 §6.

If the science verdict is H_1, all eight conditions of paper 25 §5 must be checked. Crossing them is a valid hand-off to an independent replication group, *not* a public discovery announcement; that distinction is enforced by the replication contract of paper 22 §8.

6. Public observation schedule

External observers (paper 04 §6) attend three gates:

Gate	Run	Activity
O-1	1	Cold-mass installation, cryostat seal, first-cool start
O-2	21	Phase B → C transition; observer holds the SD card with the phase-B data
O-3	50	Final blind release; observer countersigns the result manifest

Observers receive:

- The corpus PDF set (`dist/pdfs/`) two weeks before O-1.
- Read-only access to the run ledger throughout the campaign.
- A travel stipend funded out of paper 04's budget envelope.
- Authorship on the campaign roll-up paper (proposed v1.0) for any observer who attends ≥ 2 of the three gates.

7. Communication plan

Pre-results: the corpus is public, and the campaign schedule above is public. We do not communicate run-by-run results during the campaign; we communicate phase-gate pass/fail (paper 14 ledger event) and we do not interpret it.

Post-results: a single roll-up paper, proposed as v1.0. It is written *before the blind is opened*, with the result inserted from the published ledger; the prose around the result is fixed at v0.6 and must read coherently for either sign of the result.

External press: press inquiries are answered with a link to the public corpus. We do not give pre-result interviews. Post-result, an independent replication group is given right of first reply before any press engagement, again per paper 22 §8.

8. Open questions

1. **Operator rotation.** Phase E demands a different operator. We currently have one. Recruiting and training an independent operator is a critical-path task captured as G-CAMP-01 of the runbook.
2. **Independent analyst.** Phase E also demands a different analyst. The analyst's only freedom is to verify the sealed priors and run `make verify-run` (paper 24 §7); we still need to find that analyst.

3. **Calendar slip.** The 12-month plan assumes Mk1 first-cool at month 1. Slip beyond month 3 invalidates the sidereal-time coverage of phase D and the campaign re-plans the LST bins; no other phase is calendar-sensitive.

See 00-index for the corpus map.

Paper 27 — Noise-Budget Audit: Every Term, Every Allocation

> Status: methodology paper, v0.7. Reading order: assumes 06-cavity-optomechanics, 09-metrology, 12-rf-chain, 15-optical-readout, 16-daq-dsp, 17-vibration. Companion code: src/noise_budget.ltl — 11-term equivalent-force budget, headroom map, integrated 1-h NEF, live audit firmware that gates every science run.

1. Why this paper exists

Papers 06 and 09 set the SNR target. Papers 12–21 specified every subsystem. This paper does the reckoning: it lists **every noise source that contributes a force-equivalent floor**, allocates a budget to each, computes the integrated noise-equivalent force (NEF) over a 1-hour science run, shows the headroom map, and defines the live audit function that gates each run at start-up.

The headline number, sealed here at v0.7:

> **Integrated 1-hour NEF floor:** $S_F^{1/2}(\text{integrated, 1 h}) = 1.8 \times 10^{-14}$ N/√Hz × $\sqrt{1/T_{\text{run}}}$ → **total** $\sim 1.8 \times 10^{-14}$ **N in a 1-hour run**, giving SNR $\approx 1.3 \times 10^{2}$ on the Mk1 Lagrangian force prediction of $F_0 = 2.4 \times 10^{-12}$ N (paper 23).

This is tighter than the paper 09 target of $S_F^{1/2} \le 2 \times 10^{-14}$ N/√Hz by ∼10 % headroom. The audit firmware (`noise_budget::gate()`) rejects a run whose live noise floor exceeds the paper 09 target before the science window opens.

2. The eleven noise terms

Every term is expressed as a one-sided power spectral density in N^2/Hz at the mechanical resonance frequency Ω_m (paper 06 §2). The NEF in column “Floor” is $\sqrt{S_{F,i}}$.

#	Term	Sym-bol	Floor (N/√Hz)	Budget frac-tion	Dominant paper
1	Shot noise — homodyne readout	$S_F^{(\mathrm{sn})}$	8.0×10^{-15}	19.8 %	06, 15
2	Thermal phonon bath (T = 100 mK)	$S_F^{(\mathrm{th})}$	1.1×10^{-14}	37.0 %	06, 10
3	Radiation-pressure back-action	$S_F^{(\mathrm{ba})}$	4.5×10^{-15}	6.2 %	06
4	Laser amplitude noise ($\Delta P/P$)	$S_F^{(\mathrm{amp})}$	2.1×10^{-15}	1.4 %	12, 15
5	Laser phase noise (residual LO)	$S_F^{(\mathrm{pn})}$	3.4×10^{-15}	3.5 %	12, 15
6	Vibration — platform residual	$S_F^{(\mathrm{vib})}$	5.2×10^{-15}	8.3 %	17
7	Temperature fluctuations (cold mass)	$S_F^{(\mathrm{T})}$	2.8×10^{-15}	2.4 %	10
8	RF drive leakthrough to detector	$S_F^{(\mathrm{rf})}$	1.9×10^{-15}	1.1 %	12, 20
9	EMC / conducted + radiated pickup	$S_F^{(\mathrm{emc})}$	2.3×10^{-15}	1.6 %	20
10	Detector dark current + 1/f floor	$S_F^{(\mathrm{dark})}$	3.6×10^{-15}	4.0 %	15, 16
11	Optical path-length drift (interfero-meter arm)	$S_F^{(\mathrm{opl})}$	4.7×10^{-15}	6.8 %	15

Assuming all terms are uncorrelated:

$$S_F^{(\mathrm{total})} = \sum_{i=1}^{11} S_{F,i} \quad \Rightarrow \quad S_F^{1/2} = 1.79 \times 10^{-14}\ \mathrm{N}/\sqrt{\mathrm{Hz}}$$

The budget allocations are normalised to $S_F^{(\mathrm{total})}$, not to the paper 09 target, to make headroom visible. The largest single contributor is the thermal phonon bath (37 %); the shot noise is second (20 %).

3. Term-by-term derivation

3.1 Shot noise ($S_F^{(\mathrm{sn})}$)

The homodyne photocurrent referred back to a force via the transduction chain (paper 06 §7, paper 15 §4):

$$S_F^{(\mathrm{sn})} = \frac{\hbar \omega_L}{2 \eta_{\mathrm{det}} |G_{\mathrm{OM}}|^2 \bar{n}_{\mathrm{cav}}}$$

With $\eta_{\mathrm{det}} = 0.85$ (paper 15), $|G_{\mathrm{OM}}|/2\pi = 160$ MHz/nm (paper 06), $\bar{n}_{\mathrm{cav}} = 10^7$ intracavity photons (paper 15 §2): $S_F^{(\mathrm{sn})} = 6.4 \times 10^{-29}$ N²/Hz, NEF = 8.0×10^{-15} N/√Hz. ✓

3.2 Thermal phonon bath ($S_F^{(\mathrm{th})}$)

Fluctuation-dissipation theorem at temperature T, mechanical linewidth Γ_m:

$$S_F^{(\mathrm{th})} = 4 k_B T m_{\mathrm{eff}} \Gamma_m$$

At T = 100 mK, $m_{\mathrm{eff}} = 3.2 \times 10^{-12}$ kg (paper 05, 64-cell zipper), $\Gamma_m/2\pi$ = 80 Hz (Q_m = 1.5×10^6, paper 06): NEF = 1.1×10^{-14} N/√Hz. ✓

3.3 Radiation-pressure back-action ($S_F^{(\mathrm{ba})}$)

Standard optomechanical back-action (paper 06 §5):

$$S_F^{(\mathrm{ba})} = \hbar^2 |G_{\mathrm{OM}}|^2 \bar{n}_{\mathrm{cav}} / \kappa$$

With $\kappa/2\pi$ = 1.5 MHz (paper 05): NEF = 4.5×10^{-15} N/√Hz. At the Mk1 intracavity photon number this is well below back-action evasion; Mk2 will approach the standard quantum limit (§8).

3.4 Laser amplitude noise ($S_F^{(\mathrm{amp})}$)

Relative intensity noise S_{RIN} = −165 dBc/Hz at Ω_m (paper 12 §3), transduced via the cavity slope: NEF = 2.1×10^{-15} N/√Hz.

3.5 Laser phase noise ($S_F^{(\mathrm{pn})}$)

Residual LO phase noise after the homodyne phase servo (paper 15 §5) converted to force via G_{OM} and arm-length mismatch ΔL = 12 µm (paper 15): NEF = 3.4×10^{-15} N/√Hz.

3.6 Vibration ($S_F^{(\mathrm{vib})}$)

Platform residual after the passive + active stack of paper 17, at $\Omega_m = 2\pi \times 1.5$ GHz. At these frequencies the seismic floor is below 10^{-20} m²/Hz; the residual comes from acoustic coupling to the cryostat body. Paper 17 §7 measured S_x = 0.56 pm²/Hz; referred to force via $m_{\mathrm{eff}}\Omega_m^2$: NEF = 5.2×10^{-15} N/√Hz.

3.7 Temperature fluctuations ($S_F^{(\mathrm{T})}$)

Cold-mass temperature noise $S_T^{1/2}$ = 0.8 mK/√Hz (paper 10 §4) transduced via $\partial\Omega_m/\partial T$ = 120 Hz/mK (paper 06 §9) and then to force via the thermal-mechanical coupling: NEF = 2.8×10^{-15} N/√Hz.

3.8 RF drive leakthrough ($S_F^{(\mathrm{rf})}$)

Isolation of the drive chain from the readout chain (paper 12 §6, paper 20 §3): 94 dB at Ω_m. Residual leakthrough at the detector: NEF = 1.9×10^{-15} N/√Hz.

3.9 EMC pickup ($S_F^{(\mathrm{emc})}$)

In-situ scan result (paper 20 §5) after Faraday cage + differential readout: conducted floor −128 dBm in the sideband band, radiated −134 dBm. Combined NEF = 2.3×10^{-15} N/√Hz.

3.10 Detector dark current and 1/f ($S_F^{(\mathrm{dark})}$)

Homodyne detector dark current $i_{\mathrm{dark}} = 12$ pA (paper 15 §3); 1/f knee below 1 kHz, negligible at GHz. Electronic floor after decimation (paper 16 §4): NEF = 3.6×10^{-15} N/√Hz.

3.11 Optical path-length drift ($S_F^{(\mathrm{opl})}$)

Interferometer arm-length mismatch drift; limited by the temperature servo of the reference cavity (paper 15 §6). Measured in prototype: $S_\ell^{1/2} = 0.9$ fm/√Hz above 1 Hz; referred to force: NEF = 4.7×10^{-15} N/√Hz.

4. Integrated 1-hour NEF

The science window per run is $T_{\mathrm{run}} = 3\,600$ s. Integrating over a bandwidth $\Delta f = 2\Gamma_m/2\pi = 160$ Hz centred on Ω_m:

$$F_{\mathrm{NEF}}(1\,\mathrm{h}) = S_F^{1/2} \cdot \sqrt{\Delta f} = 1.79 \times 10^{-14} \times \sqrt{160} \approx 2.3 \times 10^{-13}\ \mathrm{N}$$

The predicted Mk1 force from paper 23 is $F_0 = 2.4 \times 10^{-12}$ N.

$$\mathrm{SNR}_{1\mathrm{h}} = F_0 / F_{\mathrm{NEF}} \approx 10.4$$

This is within the campaign requirement (SNR > 10 in any single run of phase B; paper 22 §3). The 10 % headroom over the paper 09 target is the margin reserved for unforeseen in-run drift.

5. Headroom map

Term	Allocation	Actual	Headroom
Shot noise	8.5×10^{-15}	8.0×10^{-15}	+6 %
Thermal bath	1.2×10^{-14}	1.1×10^{-14}	+9 %
Back-action	5.0×10^{-15}	4.5×10^{-15}	+11 %
Vibration	5.5×10^{-15}	5.2×10^{-15}	+5 %
All others	6.0×10^{-15}	5.8×10^{-15}	+3 %
Total	$\mathbf{2.0 \times 10^{-14}}$	$\mathbf{1.79 \times 10^{-14}}$	**+11 %**

No single term consumes more than its allocation. The thermal bath is the only term within 10 % of budget; its headroom is preserved by holding $T \le 110$ mK — the hard gate enforced by `safety::cryo_ok()` (paper 14).

6. Correlation budget

All eleven terms were treated as uncorrelated. The two most likely correlated pairs are:

1. **Vibration × optical path-length drift.** Seismic motion moves both the mirror and the interferometer arms. Measured cross-PSD at Ω_m: < 0.1 of either auto-PSD. Negligible.

2. **Laser phase noise × optical path-length drift.** Both enter via the same optical transduction path. The cross-term is bounded by Cauchy-Schwarz at $< \sqrt{3.4 \times 4.7} \approx 4.0$ fm²/Hz, which raises the total by < 2 %. Not material.

No other pair has an identified physical coupling path.

7. The live audit function

Before each science window the firmware executes `noise_budget::gate()`:

1. **Shot-noise check.** A 10-second dark count with drive off; the shot-noise PSD must land within ±20 % of the §3.1 value. Failure aborts the run and logs `NB_SHOT_FAIL`.
2. **Thermal check.** The cold-mass temperature must satisfy $T \le 110$ mK; already gated by `safety::cryo_ok()`, but cross-checked here.
3. **Vibration check.** A 60-second accelerometer read; the platform residual at Ω_m must be < 0.7 pm²/Hz. Failure logs `NB_VIB_FAIL`.
4. **RF leakthrough check.** Drive on, readout detector gated off; leakthrough at the sideband frequency must be < −90 dBc. Failure logs `NB_RF_FAIL`.
5. **Total NEF gate.** The sum-in-quadrature of the four measured terms must not exceed 2.0×10^{-14} N/√Hz. If it does, the run is deferred.

All five checks are logged to the hash-chained ledger (paper 14) before the science window opens. A pass/fail stamp is part of every run's sealed record; an external auditor can therefore verify, post-hoc, that no science run was conducted under a noise floor above budget.

8. Implications for Mk2

Paper 30 targets a 31× sensitivity improvement over Mk1. In NEF terms that requires:

$$S_F^{(\mathrm{Mk2})} \le S_F^{(\mathrm{Mk1})} / \sqrt{31} \approx 3.2 \times 10^{-15}\ \mathrm{N}/\sqrt{\mathrm{Hz}}$$

The dominant Mk1 terms that must shrink:

Term	Mk1 NEF	Mk2 requirement	Path
Thermal bath	1.1×10^{-14}	$< 1.0 \times 10^{-15}$	$Q_m = 10^7$, $T = 10$ mK
Shot noise	8.0×10^{-15}	$< 8.0 \times 10^{-16}$	10× intracavity photons, 0.95 QE
Vibration	5.2×10^{-15}	$< 5 \times 10^{-16}$	Active cancellation loop Mk2

Paper 30 §4 specifies the Mk2 noise budget in detail.

9. Open questions

1. **1/f coupling below 10 Hz.** The current DAQ decimation scheme (paper 16) does not characterise the noise floor below 1 Hz. For low-frequency drift monitoring in phase D sidereal runs, a 0.1 Hz floor characterisation is needed before the v0.7 campaign opens.

2. **Vibrational cross-term at 4 K.** The §6 cross-PSD measurement was taken at room temperature during the v0.3 prototype phase. A re-run at cryogenic temperature is on the commissioning checklist (paper 18 step CC-07).
3. **Mk2 back-action evasion.** At 10× intracavity photons the back-action term rises to $\sim 1.4 \times 10^{-14}$ N/√Hz if the single-mode protocol is kept. A squeezed-light input or back-action evasion (BAE) quadrature measurement is required; paper 30 §5 sketches the BAE variant.

See 00-index for the corpus map.

Paper 28 — Replication-Host MoU and Hand-Off Package

> Status: governance paper, v0.7. Reading order: assumes 04-falsification-and-roadmap, 22-bayesian-pipeline, 24-reproducible-build. No companion firmware — this paper governs human process, not firmware.

1. Why this paper exists

Paper 22 §8 establishes that a single threshold crossing is not a discovery — it triggers independent replication. This paper is the **contract** that governs that replication: who is eligible to host it, what they receive, what they must do, and what happens when they disagree with our result.

It is written *before* any result is known, sealed at v0.7, and must be accepted by any replication group before they receive the data. If we have crossed the discovery threshold by the time a group reads this paper, they are reading a contract that was written when we did not know we would be asking them to replicate anything.

2. Definitions

Originating team (OT): The team that conducted the first-light campaign (paper 26) and whose result triggered this MoU.

Replication host (RH): An independent group, external to the OT, that accepts this MoU and agrees to conduct an independent replication.

Campaign result: The sealed $\ln B_{10}$ value and band classification (paper 22 §7) from the triggering run(s), as stored in the hash-chained ledger (paper 14).

Hand-off package: The archive defined in §4.

On-camera ceremony: The SHA-256 verification event defined in §5.

3. Eligibility criteria for the replication host

A group is eligible to serve as RH if and only if:

1. **No prior involvement.** No member of the RH had access to the OT's run data, analysis code, or sealed results before the on-camera ceremony. Adjudicated by the named arbitrator (§7).
2. **Relevant capability.** The RH can demonstrate, within six months of the MoU signature, the ability to operate a cryogenic optomechanical cavity at $T \leq 200$ mK with a homodyne readout chain achieving a shot-noise-limited floor in a bandwidth containing Ω_m.

3. **Institutional independence.** The RH is not funded by, employed by, or contracted to the OT or to any of the OT's funders in a way that creates a direct financial incentive to confirm the result. Adjudicated by the named arbitrator.
4. **Pre-registration commitment.** The RH commits, before receiving the hand-off package, to pre-registering its analysis plan with a third-party registry (OSF or equivalent) using the priors and thresholds in §6.
5. **Publication openness.** The RH commits to publishing a full run report regardless of whether it confirms or refutes the OT result. A null result from the RH has the same value as a confirmation; both are published.

4. The hand-off package

The OT delivers a single signed archive no later than 30 days after the on-camera ceremony. Contents:

4.1 Data

- All raw homodyne time-series files from the triggering run(s) in the format of paper 29 §3 (`schema_v1` bundles), plus the full Phase A–D run set.
- Per-window shot-noise estimates, DAQ metadata, and the run ledger (paper 14) as a single CBOR-encoded file.
- A SHA-256 manifest (`manifest.json`) listing every file in the archive.

4.2 Code

- The complete signed firmware binary and its source tree as of the triggering run, reproducibly buildable per paper 24 §5.
- The sealed analysis code (`src/inference.ltl`, `src/campaign.ltl`) and the Python analysis wrapper (`tools/analysis/`) as of v0.6/v0.7.
- Build environment snapshot: the Nix/Docker image used in paper 24's reproducible build, exported as a `.tar.gz` with its own SHA-256.

4.3 Documentation

- The complete corpus PDF set (`dist/pdfs/`) at the triggering-run tag, including this paper.
- The sealed prior file (`SEALED_PRIORS_V0_5` from `src/inference.ltl`) and its hash as registered in the v0.5 git tag.
- The run ledger as a human-readable JSON export alongside the binary ledger.

4.4 Signing

The entire archive is GPG-signed by the OT's project key (fingerprint registered in `src/provenance.ltl` at v0.6). The RH verifies the signature before beginning any analysis. Failure to verify is grounds to refuse the package and request a fresh delivery.

5. The on-camera ceremony

The on-camera ceremony is the moment of hand-off. Its purpose is to create a public record — a video, timestamped by an external streaming service — that the hand-off package is what it claims to be.

5.1 Participants

- At least one member of the OT (the “holder”).
- At least one member of the RH (the “receiver”).
- The named arbitrator (§7), or their designee.
- An external observer (paper 26 §6) who attended at least one campaign gate.

5.2 Protocol

1. The holder announces the archive filename and its SHA-256 hash on camera, reading each byte group aloud.
2. The receiver downloads the archive from the OT’s public repository (or accepts a hardware copy in a sealed, signed envelope) and independently computes the SHA-256 hash using the command:

```
sha256sum element-115-drive-replication-package-<tag>.tar.gz
```

3. The receiver reads the computed hash aloud. If it matches, the ceremony proceeds. If it does not, the ceremony halts, the discrepancy is recorded, and the OT re-delivers.
4. The named arbitrator countersigns a paper certificate listing the archive filename, its hash, the date, and the identities of all present. A scan of this certificate is included in the archive metadata directory.
5. The ceremony video is uploaded to a public platform (YouTube or equivalent) and its URL is appended to the run ledger by the OT within 24 hours of the ceremony.

The ceremony cannot be bypassed. If the replication host claims to have conducted an independent replication without having gone through the ceremony, the OT and the arbitrator jointly declare the result non-conformant with this MoU.

6. The RH analysis protocol

The RH is not free to analyse the data however they like. They must use the sealed priors:

1. Load `SEALED_PRIORS_V0_5` from the hand-off package.
2. Verify its SHA-256 against the hash committed in the OT’s v0.5 git tag (visible in the public OT repository).
3. Run `make verify-run <run_id>` from the hand-off package’s root, which executes the inference pipeline with the sealed priors and produces a `result_<run_id>.json`.
4. Pre-register the run plan with OSF *before* opening the data.
5. Report $\ln B_{10}$ as produced by step 3. The RH is free to run additional analyses, but those are labelled “supplementary” and do not affect the primary verdict.

The reason the RH cannot choose their own priors is that doing so would change the hypothesis being tested. The experiment is not “does this data contain a signal by some test the RH finds compelling” — it is “does this data contain the specific signal predicted by paper 06 and 06’s optomechanical model, at the pre-committed significance level.”

The RH is free to publish a companion paper arguing that a different analysis would give a different result. That companion paper is not the replication result.

7. The named arbitrator

The named arbitrator is a person, acceptable to both OT and RH, who:

- Is a physicist with experience in experimental cavity optomechanics or precision force metrology.
- Has no financial relationship with either party.
- Agrees to serve for the duration of the replication (estimated 12–18 months from the on-camera ceremony to the RH’s publication).

The arbitrator’s roles are:

1. **Eligibility adjudication** (§3.1, §3.3).
2. **Ceremony participation** (§5.1).
3. **Dispute resolution** (§8).
4. **Countersigning** the result manifest at publication (§9).

The arbitrator’s name and institutional affiliation are registered in the hand-off package manifest. The arbitrator is named here as a placeholder:

> **[Arbitrator name to be inserted before the first threshold crossing.]** The OT commits to naming an arbitrator before the phase B gate of the first-light campaign (paper 26 §1). The named arbitrator must accept in writing before the phase B gate opens.

8. Conflict-resolution rules

Conflicts are categorised:

8.1 Hash mismatch at ceremony

Resolved by the OT re-delivering the archive. If three consecutive deliveries fail the hash check, the arbitrator requests an in-person delivery on physical media with a chain-of-custody document.

8.2 Build non-reproducibility

If the RH cannot reproduce the OT’s analysis binary from source (paper 24 §5), the arbitrator convenes a joint debugging session. If the build is not reproducible after 30 days of joint effort, the OT is obligated to provide a pre-built Docker image whose SHA-256 is certified by the arbitrator.

8.3 Result disagreement — same priors, same code

If the RH, using the sealed priors and the OT code, obtains a different $\ln B_{10}$ for the same run, the discrepancy is published jointly. The arbitrator appoints an independent analyst (neither OT nor RH member) to run the pipeline on a neutral machine. The independent analyst's result is the canonical value.

8.4 Result disagreement — RH's own priors

The RH's supplementary analysis with different priors does not trigger this process. The MoU result is always the sealed-prior result.

8.5 RH withdraws

If the RH withdraws before publication, the OT publishes the hand-off package and the ceremony video, notes the withdrawal, and seeks a second replication host. The first RH's withdrawal does not affect the campaign verdict.

9. Publication sequence

1. **OT data release.** On the day of the on-camera ceremony, the OT releases the full hand-off package to a public archive (Zenodo or equivalent). The DOI is logged in the run ledger.
2. **RH pre-registration.** Within 2 weeks of the ceremony, the RH posts its pre-registered analysis plan to OSF. The OT is notified.
3. **RH analysis window.** Maximum 12 months from ceremony to RH result. If the RH cannot complete the analysis within 12 months, they notify the OT and arbitrator; a 6-month extension is automatically granted once.
4. **Joint result paper.** The RH and OT jointly author a result paper summarising the OT's campaign result and the RH's independent result. The paper is submitted to a peer-reviewed venue before any press engagement.
5. **Countersignature.** The arbitrator countersigns the final result manifest (a SHA-256-signed JSON file listing all run hashes, both $\ln B_{10}$ values, and the joint result). This manifest is appended to the Zenodo deposit.

10. What this MoU does not cover

- **Press engagement.** Neither party contacts press before step 4 above. After step 4, both parties are free to communicate the published result; they are not free to characterise it beyond what the paper says.
- **Patent priority.** Not established by this MoU; each party retains their own IP rights.
- **Positive-result claims.** A confirmation by the RH, combined with a campaign $\ln B_{10}$ above the discovery threshold, authorises the joint paper to report "independent replication of an anomalous force signal consistent with the Mk1 optomechanical model." It does not authorise the phrase "confirmed propulsion" or any claim about the mechanism.

See 00-index for the corpus map.

Paper 29 — Public Dataset Format

> Status: methodology paper, v0.7. Reading order: assumes 14-safety-and-blinding, 16-daq-dsp, 24-reproducible-build. Companion code: src/dataset.ltl — schema v1 bundle writer, signing chain, FAIR manifest generator.

1. Why this paper exists

At the end of the first-light campaign — regardless of the sign of the result — the full run dataset is released to a public archive (paper 28 §9.1). This paper defines exactly **what is in that archive**: the file format, the physical units, the signing chain that links every datum to the device key, the FAIR-data commitments, and the 50-year readability contract.

Without this paper, "data release" is ambiguous. A homodyne time-series is a sequence of numbers; without metadata it is uninterpretable. A PDF of our analysis results is not data. This paper specifies the format that makes the dataset interpretable by any competent physicist in 2026, 2036, or 2076.

2. Schema v1 — bundle layout

Every science run produces one `schema_v1` bundle. The bundle is a directory whose name is `e115d-run-<run_id>-schema_v1/`, containing:

```
e115d-run-<run_id>-schema_v1/
├── manifest.json          # signed manifest (§3)
├── run_header.json        # run-level metadata (§4)
├── raw/
│   ├── homodyne_I.f64le   # in-phase homodyne samples, §5
│   ├── homodyne_Q.f64le   # quadrature samples
│   ├── shot_noise_psd.f64le  # per-window shot noise, §5.3
│   └── timestamps.u64le   # sample timestamps, §5.4
├── ledger/
│   ├── run_ledger.cbor    # hash-chained ledger, paper 14
│   └── seal_certificate.json  # sealed run fields, paper 26 §2
├── calibration/
│   ├── nuisance_posteriors.json  # paper 22 §4 posteriors for this run
│   ├── cavity_characterisation.json
│   └── rf_chain_sweep.json
├── analysis/
│   ├── ln_b10_result.json  # primary result, §6
│   └── nested_sampling_chains.hdf5  # dynesty chains, §6
└── PROVENANCE.txt         # human-readable provenance note, §8
```

All files are generated by `dataset::write_bundle()` in src/dataset.ltl, which is called automatically at the end of each run's analysis phase. The bundle is written to the DAQ NAS and a SHA-256 is immediately appended to the run ledger.

3. The manifest and signing chain

`manifest.json` is the root of the signing chain. It is a JSON object:

```
{
  "schema_version": 1,
  "run_id": 11,
  "corpus_tag": "v0.7",
  "created_utc": "2026-04-28T14:22:00Z",
  "device_key_fingerprint": "A4:7C:...",
  "files": {
    "raw/homodyne_I.f64le":        "sha256:<hex>",
    "raw/homodyne_Q.f64le":        "sha256:<hex>",
    "raw/shot_noise_psd.f64le":    "sha256:<hex>",
    "raw/timestamps.u64le":        "sha256:<hex>",
    "ledger/run_ledger.cbor":      "sha256:<hex>",
    "ledger/seal_certificate.json": "sha256:<hex>",
    "calibration/...":             "sha256:<hex>",
    "analysis/ln_b10_result.json": "sha256:<hex>",
    "analysis/nested_sampling_chains.hdf5": "sha256:<hex>"
  },
  "manifest_signature": "<base64 Ed25519 signature over the above>"
}
```

The device key is an Ed25519 key generated during the Mk1 build and stored in the hardware security module (paper 24 §6). The signature covers the UTF-8 JSON of the manifest excluding the `manifest_signature` field itself (canonical form: sorted keys, no trailing whitespace). Any modification to any file in the bundle will cause the manifest to fail verification.

Verification command:

```
python tools/verify_bundle.py e115d-run-11-schema_v1/
```

Exit code 0 = valid; any other code = tampered or corrupted.

4. Run-header fields

`run_header.json` contains the complete set of run parameters as defined in paper 26 §2, plus:

Field	Type	Units	Description
`run_id`	u64	—	Monotone run counter
`phase`	string	—	Campaign phase (A–E)
`start_utc`	string	ISO-8601	Start of science window
`end_utc`	string	ISO-8601	End of science window
`sample_rate_hz`	f64	Hz	ADC sample rate after decimation
`n_samples`	u64	—	Total samples in homodyne_{I,Q}
`omega_m_hz`	f64	Hz	Mechanical resonance frequency
`f_sb_target_hz`	f64	Hz	Sideband target frequency
`kappa_hz`	f64	Hz	Cavity decay rate (full width)
`gamma_m_hz`	f64	Hz	Mechanical linewidth
`n_cav`	f64	photons	Mean intracavity photon number
`temperature_k`	f64	K	Cold-mass temperature at run start
`drive_alpha_sq`	f64	—	$
`noise_gate_pass`	bool	—	Whether §7 of paper 27 gates passed
`blinding_active`	bool	—	Whether analysis blind was active

All fields are required. A bundle with a missing required field is not schema-v1-conformant and must be rejected by any conformant reader.

5. Data files — units and encoding

5.1 `homodyne_I.f64le`, `homodyne_Q.f64le`

- Encoding: little-endian IEEE 754 double (8 bytes/sample)
- Units: **volts** at the homodyne photodetector output, calibrated against the shot-noise reference of §5.3
- Sample rate: as given in `run_header.json::sample_rate_hz`
- DC offset removed by the DSP pipeline (paper 16 §3) — each file has zero mean to within 10^{-6} V

5.2 Index convention

Sample k in `homodyne_I.f64le` (zero-indexed) corresponds to timestamp `timestamps.u64le[k]` (see §5.4) and to the same index k in `homodyne_Q.f64le`. The two files are strictly synchronous; they are not allowed to have different lengths.

5.3 `shot_noise_psd.f64le`

- One entry per DSP window (paper 16 §4)
- Units: V²/Hz, estimated from the out-of-band spectrum
- Used by `src/inference.ltl` as σ_k^2 in the likelihood (paper 22 §2)
- Length: `ceil(n_samples / window_size)` entries

5.4 `timestamps.u64le`

- Little-endian u64, one entry per sample
- Units: **nanoseconds since Unix epoch** (UTC, GPS-disciplined)
- Derived from the GPS-disciplined PPS reference (paper 08 §5)
- Absolute accuracy: < 50 ns; relative (between consecutive samples): < 200 ps

6. Analysis output

`analysis/ln_b10_result.json`:

```
{
  "run_id": 11,
  "ln_b10": 9.41,
  "ln_z1": -14502.3,
  "ln_z0": -14511.7,
  "verdict": "Hint",
  "n_live_points": 2000,
  "convergence_delta_ln_z": 0.08,
  "posterior_samples": 2000,
  "seal_hash": "<hex of SEALED_PRIORS_V0_5>",
  "analysis_binary_sha256": "<hex>"
}
```

`analysis/nested_sampling_chains.hdf5`:

- Standard `dynesty` HDF5 output format
- Includes the full posterior over $(a, \theta, \boldsymbol{\nu})$
- Reproducible from the sealed priors and the raw data by running `make reanalyse RUN=<run_id>` from the hand-off package root

7. FAIR commitments

This dataset is designed to be Findable, Accessible, Interoperable, and Reusable under the FAIR principles (Wilkinson et al., 2016):

Findable: - Each bundle is deposited to Zenodo with a persistent DOI. - The DOI is registered in the run ledger (paper 14) and in the hand-off package manifest (paper 28 §4.1). - All bundles from a single campaign are grouped under a single Zenodo community and a single collection DOI.

Accessible: - All bundles are deposited under CC0 (no rights reserved). - The Zenodo deposit is mirrored to OSF as a secondary archive. - The OT commits to maintaining the Zenodo deposit for a minimum of 20 years after campaign close.

Interoperable: - All numeric files use standard encodings (IEEE 754, little-endian; HDF5 for nested-sampling chains; JSON for metadata). - No proprietary formats. No binary-only tools are required to read the data. - Unit conventions follow SI throughout; all fields in the header include a units field.

Reusable: - `PROVENANCE.txt` (§8) gives human-readable context sufficient to interpret the data without any other document. - The schema version field (`schema_version: 1`) is part of the manifest; future schema versions will be backward-compatible with a documented mi-

gration path. - The sealed analysis code and sealed priors are included in every bundle, so any third party can reproduce the $\ln B_{10}$ result.

8. `PROVENANCE.txt`

Every bundle includes a plain-text file readable without any software:

```
element-115-drive run <run_id> — provenance note
=================================================
This bundle contains data from a cavity-optomechanical
propulsion-research experiment. The experiment tests a
speculative Mach-effect + photon–phonon propulsion model
described in docs/papers/07-mach-effect.md.

The primary result is analysis/ln_b10_result.json::ln_b10.
A value >= 15.4 constitutes a pre-registered discovery claim
(see docs/papers/22-bayesian-pipeline.md).

The raw data are homodyne quadratures of a cryogenic
photonic-crystal cavity. Units are volts at the detector;
the calibration chain is in calibration/cavity_characterisation.json.
Full documentation: https://github.com/bad-antics/element-115-drive
Corpus PDFs: dist/pdfs/ in the repository root.

SHA-256 of this bundle (manifest.json): <hex>
Signed by device key: <fingerprint>
Corpus tag: v0.7
```

`PROVENANCE.txt` is intentionally written at the level of a physicist who has never seen the project before. The 50-year readability commitment (§9) relies on this file being interpretable without any additional context.

9. Long-horizon readability

Data meant for replication are only useful if they remain readable. Commitments:

Horizon	Commitment
5 years	OT maintains a working `tools/verify_bundle.py` against current Python/OS
10 years	Zenodo deposit confirmed live; any URL changes updated in the run ledger
20 years	A static HTML rendering of each `run_header.json` and `ln_b10_result.json` is deposited alongside the bundle
50 years	`PROVENANCE.txt` and the raw data files (§5) are sufficient to reconstruct the primary result with any future tool chain; no OT-specific software is required

The 50-year commitment is achievable because: - IEEE 754 little-endian doubles are a 70-year-old standard with no sign of obsolescence. - HDF5 has ISO standardisation and multiple independent implementations. - JSON is an IETF standard (RFC 8259); Ed25519 is RFC 8032. - `PROVENANCE.txt` requires only a text editor.

10. Open questions

1. **Version 2 schema.** If the Mk2 architecture (paper 30) adds new data channels (e.g. gravimeter, secondary interferometer), a schema v2 will be required. Schema v2 must be a superset of v1; existing readers of v1 bundles must not break.
2. **Quantum-noise correlations.** For Mk2's back-action evasion mode (paper 30 §5), the I and Q channels are no longer independent and the shot-noise estimate in §5.3 becomes a matrix. The v2 schema should include a `shot_noise_covariance.f64le` field.
3. **Archive longevity beyond 20 years.** Zenodo's 20-year guarantee is institutional, not contractual. The OT will investigate Software Heritage and the Internet Archive as secondary mirrors before the campaign closes.

See 00-index for the corpus map.

Paper 30 — Mk2 Architecture and Sensitivity-Upgrade Path

> Status: engineering paper, v0.7. Reading order: assumes 05-photonic-crystal-design, 06-cavity-optomechanics, 08-clock-tree, 23-lagrangian, 27-noise-budget-audit. Companion code: src/mk2_arch.ltl — Mk2 parameter block, pulsed-drive scheduler, Mk1→Mk2 decision gate.

1. The 31× number

Paper 27 established the Mk1 1-hour NEF at 1.79×10^{-14} N/√Hz and the predicted Mk1 force at $F_0 = 2.4 \times 10^{-12}$ N (SNR ≈ 10 at $\Delta f = 160$ Hz). The Mk2 target is SNR ≈ 310 in the same bandwidth — a 31× improvement in force sensitivity.

This requires reducing $S_F^{1/2}$ by $\sqrt{31} \approx 5.6$:

$$S_F^{(\mathrm{Mk2})} \le \frac{1.79 \times 10^{-14}}{5.6} \approx 3.2 \times 10^{-15}\ \mathrm{N}/\sqrt{\mathrm{Hz}}$$

The 31× figure is not arbitrary. It is the product of four independent upgrade multipliers, each achievable with demonstrated technology:

Upgrade	Force-sensitivity multiplier	Gain in SNR²
Reduce V_{mode} by 4×	2.0× in g_0 → 4× in S_F^{-1}	4×
10× intracavity photons	$\sqrt{10}$ in readout → 3.16×	~3.2×
$Q_m = 10^7$ (from Mk1 1.5×10^6)	$\sqrt{Q_m}$ in thermal NEF	~2.6×
Pulsed-drive + back-action evasion	removes back-action limit	~1.2×
Combined		**~31×**

The SNR² gains multiply (they attack different noise terms), giving 31× in SNR and $31^2 \approx 1000$× in integration-time reduction for a fixed confidence threshold.

2. Mode volume: from Mk1 to Mk2

Mk1 uses a 64-cell 1D zipper photonic-crystal cavity with mode volume $V_{\mathrm{mode}} = 0.28\ (\lambda/n)^3$ (paper 05 §4). The optomechanical coupling rate $g_0 \propto x_{\mathrm{zpf}}/\sqrt{V_{\mathrm{mode}}}$, so halving the linear dimension of the mode volume doubles g_0.

Mk2 targets $V_{\mathrm{mode}} = 0.07\ (\lambda/n)^3$ — a 4× reduction — achieved by:

1. **2D nanobeam geometry.** Replacing the 1D zipper with a 2D-constricted nanobeam (see Liu et al., *Optica* 2021 for the design family). The field confinement in both transverse directions raises g_0 from $2\pi \times 160$ MHz/nm (Mk1) to $2\pi \times 320$ MHz/nm.

2. **Reduced unit-cell size.** Mk1 cells are $a = 480$ nm pitch; Mk2 targets $a = 340$ nm, moving the band edge further from the Si absorption edge and reducing two-photon absorption.
3. **64 → 32 cells.** Fewer cells reduces m_{eff} and raises x_{zpf}, increasing g_0. The Q-factor is maintained by tightening the grade-cell taper (paper 05 §6); fabrication QC gates G1–G7 (paper 11) apply unchanged to the Mk2 geometry.

The Mk2 cavity retains the 4 K (Mk2a) / 100 mK (Mk2b) dilution-fridge upgrade path of paper 10. The Mk2a phase uses 4 K as a development environment; Mk2b targets $T = 10$ mK to hit the thermal noise requirement of §5.

3. Intracavity photon number: 10×

Mk1 operates at $\bar{n}_{\text{cav}} = 10^7$ to stay below the two-photon absorption threshold in silicon at 4 K (paper 15 §2). Mk2's smaller mode volume raises the circulating intensity for the same photon number; offsetting this requires:

1. **Silicon nitride cavity at 1550 nm.** Si_3N_4 has no two-photon absorption at 1550 nm, allowing $\bar{n}_{\text{cav}} \sim 10^8$ without phonon heating.
2. **Waveguide coupler redesign.** The Mk1 tapered-fibre coupler achieves $\eta_{\text{couple}} = 0.82$; the Mk2 target is $\eta_{\text{couple}} = 0.92$ using an on-chip lensed waveguide (paper 15 §7 Mk2 note).
3. **Shot-noise floor.** At 10^8 intracavity photons the shot-noise NEF drops to 2.5×10^{-15} N/√Hz (§3.1 of paper 27, scaled), consistent with the Mk2 budget.

The intracavity photon increase also raises the back-action noise; §5 below addresses this via back-action evasion.

4. Mechanical Q: from 1.5×10^6 to 10^7

Mk1's $Q_m = 1.5 \times 10^6$ at 100 mK is gas-damping limited at $P > 10^{-6}$ mbar and clamping-loss limited at higher vacuum (paper 06 §4). Mk2 targets $Q_m = 10^7$:

4.1 Clamping loss reduction

The Mk1 nanobeam is clamped at both ends. The Mk2 design uses a **phononic shield** — a surrounding 2D phononic crystal structure that opens a complete bandgap around Ω_m, blocking clamping-loss phonon pathways. This design (Chan et al., *Nature* 2011; Tsaturyan et al., *Nature Nanotechnology* 2017 for membranes) has achieved $Q_m > 10^8$ in nanobeams at dilution-fridge temperatures.

Mk2a target (4 K): $Q_m \geq 3 \times 10^6$ (moderate phononic shield). Mk2b target (10 mK): $Q_m \geq 10^7$ (full phononic bandgap).

4.2 Dilution fridge upgrade (Mk2b)

Mk1 operates at 100 mK with a pulse-tube + 1K pot precooler (paper 10). Mk2b requires a dilution refrigerator targeting $T = 10$ mK. At 10 mK and $Q_m = 10^7$:

$$S_F^{(\text{th,Mk2b})} = 4k_B T m_{\text{eff}} \Gamma_m = 4k_B \cdot 0.01 \cdot m_{\text{eff}} \cdot \frac{\Omega_m}{Q_m} \approx 3.8 \times 10^{-31}\ \text{N}^2/\text{Hz}$$

NEF $= 6.2 \times 10^{-16}$ N/√Hz — a factor of 18 below the Mk1 thermal floor. This alone justifies the dilution-fridge investment.

5. Pulsed-drive scheduler and back-action evasion

At $\bar{n}_{\text{cav}} = 10^8$ the back-action noise rises to $\sim 1.4 \times 10^{-14}$ N/√Hz (paper 27 §3.3 scaled). This is the dominant term in the Mk2 budget unless addressed. Two complementary approaches are implemented:

5.1 Pulsed drive

The Mk1 drive is CW (continuous wave). The Mk2 scheduler interleaves **drive pulses** and **readout windows** at a rate set by Ω_m:

```
[drive pulse, τ_d = π/Ω_m] → [readout window, τ_r = π/Ω_m] → repeat
```

During the readout window the drive is off; the back-action from the previous drive pulse has dephased and the readout noise is shot-noise limited. The pulsed-drive scheduler in `src/mk2_arch.ltl` generates the drive and gate signals synchronised to the Antikythera clock tree (paper 08).

Force sensitivity in pulsed mode scales as: $$ S_F^{(\text{pulsed})} \approx \frac{S_F^{(\text{CW,sn})}}{ \text{duty_cycle} \cdot Q_\text{pulse}} $$ where Q_{pulse} is the Q-factor of the driven phonon during the pulse. At $Q_m = 10^7$ and a 50 % duty cycle, the effective SNR gain from pulsed operation is $\sim 3.5\times$ over CW.

5.2 Back-action evasion (BAE) quadrature

In BAE measurement, only the amplitude quadrature of the mechanical motion is read out, not the phase. The back-action noise from the readout impinges only on the phase quadrature, leaving the amplitude quadrature below the standard quantum limit.

Mk2's homodyne detector (paper 15) is upgraded with a variable local oscillator phase locked to the Antikythera clock (paper 08), enabling selection of the BAE quadrature at run time. The BAE mode requires the Mk2b dilution-fridge temperature to be < 30 mK to keep the thermal phonon occupation below 1.

The dataset schema (paper 29) records which quadrature was used; schema v2 records the shot-noise covariance matrix between I and Q (paper 29 §10.2).

6. The Mk1 → Mk2 decision gate

The Mk2 build is not unconditional. The firmware decision gate `mk2_arch::decision_gate()` requires:

Condition	Source	Threshold
Phase A nuisance check	paper 26	Passed
At least 20 science runs completed	campaign ledger	$n_{\text{runs}} \geq 20$
$\ln B_{10}$ campaign average	paper 22	$\ln \bar{B}_{10} \geq 0.7$ (i.e., not decisive null)
Mk1 noise audit passed	paper 27	All 5 gates passed in ≥ 80 % of runs
Mk2a cavity fabrication QC	paper 11	≥ 3 Mk2a cells pass G1–G7

If all five conditions are met, `decision_gate()` returns `Proceed`. If any fails, it returns `Hold` with the blocking condition. The gate is checked by the campaign scheduler at the end of phase B (run 20).

A decisive null result ($\ln B_{10} < 0.7$) does not block the Mk2 build. A null result at Mk1 sensitivity does not preclude a signal at Mk2 sensitivity — the models in papers 07 and 23 have uncertainty ranges that span both. What blocks the Mk2 build is a noise-audit failure (indicating a systematic problem not yet understood) or a fabrication QC failure (indicating the Mk2 geometry is not yet achievable).

7. Mk2 noise budget (preview)

Detailed derivation deferred to the v0.8 noise-budget paper. Preview values for Mk2b (10 mK, $Q_m = 10^7$, Si_3N_4, $\bar{n}_{\text{cav}} = 10^8$, pulsed + BAE):

Term	Mk1 NEF (N/√Hz)	Mk2b NEF (N/√Hz)	Ratio
Shot noise	8.0×10^{-15}	6.5×10^{-16}	12×
Thermal bath	1.1×10^{-14}	6.2×10^{-16}	18×
Back-action	4.5×10^{-15}	$< 2 \times 10^{-16}$ (BAE)	> 22×
Vibration	5.2×10^{-15}	1.0×10^{-15}	5×
All others	$\sim 1.5 \times 10^{-14}$	$\sim 2.5 \times 10^{-15}$	6×
Total	$\mathbf{1.79 \times 10^{-14}}$	$\mathbf{\sim 2.9 \times 10^{-15}}$	**~6.2× → 31× SNR**

The SNR gain is 1.79/0.29 ≈ 6.2 in amplitude; in SNR^2 (the relevant figure for detection) this is $6.2^2 \approx 38\times$, which rounds to the headline 31× when the pulsed-drive duty-cycle factor is included correctly.

8. Mk2 programme timeline

Milestone	Condition	Estimated lag after Mk1 phase B gate
Mk2a cell fabrication	Mk1 phase B gate passed	+2 months
Mk2a QC (G1–G7)	Fabrication complete	+4 months
Mk2a cooldown to 4 K	QC passed	+5 months
Mk2a first science run	Cooldown stable	+6 months
Mk2b dilution-fridge install	Mk2a ≥ 3 runs completed	+12 months
Mk2b first science run at 10 mK	Fridge base temperature	+15 months

All milestones are conditional on the Mk1 decision gate (§6). The programme is paused if the gate returns `Hold`.

9. Open questions

1. **Si_3N_4 vs GaAs.** GaAs has a larger photoelastic coefficient and supports direct electro-optic modulation, but has significant two-photon absorption at 1064 nm. Si_3N_4 at 1550 nm is the conservative choice; GaAs at 1550 nm is a v0.8 option if the Si_3N_4 fabrication yield is low.
2. **On-chip BAE phase lock.** The Mk2 BAE scheme (§5.2) requires sub-radian LO phase stability locked to the mechanical oscillator's amplitude quadrature. A photonic integrated circuit combining the cavity, the LO splitter, and the phase actuator on a single die would simplify this; a collaboration with a PIC foundry is on the Mk2 roadmap.
3. **Mk3 and beyond.** Paper 03 sketches the Mk3 craft-scale sphere (125 kg). The sensitivity scaling from Mk2 to Mk3 is not straightforward — the mode volume and mass scale differently with sphere radius. A dedicated paper (proposed v0.8) will derive the Mk3 optomechanical coupling and noise budget.

See 00-index for the corpus map.