

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 $\sqrt{31} \approx 5.6$:

$$S_F^{(\text{Mk2})} \leq \frac{1.79 \times 10^{-14}}{5.6} \approx 3.2 \times 10^{-15} \text{ N}/\sqrt{\text{Hz}}$$

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 SNR ²
Reduce V_{mode} by 4×	$2.0\times$ in $g_0 \rightarrow 4\times$ in S_F^{-1}	4×
10× intracavity photons	$\sqrt{10}$ in readout $\rightarrow 3.16\times$	$\sim 3.2\times$
$Q_m = 10^7$ (from Mk1 1.5×10^6)	$\sqrt{Q_m}$ in thermal NEF	$\sim 2.6\times$
Pulsed-drive + back-action evasion	removes back-action limit	$\sim 1.2\times$
Combined		$\sim 31\times$

The SNR² gains multiply (they attack different noise terms), giving 31× in SNR and $31^2 \approx 1000\times$ 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 $V_{\text{mode}} = 0.28 (\lambda/n)^3$ (paper 05 §4). The optomechanical coupling rate $g_0 \propto \kappa_{\text{zpf}}/\sqrt{V_{\text{mode}}}$, so halving the linear dimension of the mode volume doubles g_0 .

Mk2 targets $V_{\text{mode}} = 0.07 (\lambda/n)^3$ — a 4× reduction — achieved by:

- 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 χ_{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 $n_{\text{cav}} = 10^7$ 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 $n_{\text{cav}} \sim 10^8$ 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 $Q_m = 10^7$:

$$S_F^{(\text{th}, \text{Mk2b})} = 4k_B T m_{\text{eff}} \Gamma_m = 4k_B \cdot 0.01 \cdot m_{\text{eff}} \cdot \frac{\Omega_m}{Q_m} \approx 3.8 \times 10^{-31} \text{ N}^2/\text{Hz}$$

NEF = 6.2×10^{-16} N/ $\sqrt{\text{Hz}}$ — 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/ $\sqrt{\text{Hz}}$ (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,  $\tau_d = \pi/\Omega_m$ ] → [readout window,  $\tau_r = \pi/\Omega_m$ ] → 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 $\geq 80\%$ 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/ $\sqrt{\text{Hz}}$)	Mk2b NEF (N/ $\sqrt{\text{Hz}}$)	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)	$> 22\times$
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	1.79×10^{-14}	$\sim 2.9 \times 10^{-15}$	$\sim 6.2\times \rightarrow 31\times$ SNR

The SNR gain is $1.79/0.29 \approx 6.2$ in amplitude; in SNR^2 (the relevant figure for detection) this is $6.2^2 \approx 38\times$, which rounds to the headline $31\times$ 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 \geq 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

- Si₃N₄ vs GaAs.** GaAs has a larger photoelastic coefficient and supports direct electro-optic modulation, but has significant two-photon absorption at 1064 nm. Si₃N₄ at 1550 nm is the conservative choice; GaAs at 1550 nm is a v0.8 option if the Si₃N₄ fabrication yield is low.
- 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.
- 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.

element-115-drive · 30 · v0.6 · lateralus.dev/research/element-115-drive (<https://lateralus.dev/research/element-115-drive>)