Structured estimation of thermal drift in WDM microring banks

How much resonance stability is lost to the cadence of the monitor, and how much is recovered without touching the hardware. Reproducible simulation on public parameters only.

1 · SCOPE

This note quantifies a specific and avoidable loss of margin in WDM microring banks: the one that accumulates **between monitor probes**. Every parameter comes from a material constant or the open literature, and each is listed with its provenance in Table 1. **No foundry-proprietary data was used.** The estimator evaluated here is its generic form — a linear filter with an explicit state structure; the production engine is not reproduced. The purpose is methodological: to show the problem is tractable in software before discussing anyone's measured data.

2 · THE PROBLEM

A WDM microring bank shares a substrate. The resonance monitor **cannot probe every channel at once**: the dither needed to estimate detuning degrades the channel in service, so the monitor runs *round-robin*, one channel per dwell period. At **16 channels and 60 s dwell, each ring receives its own measurement only once every 16 minutes**. Between its own probes, conventional control freezes its last correction and runs blind.

That blind interval is where drift accumulates. And the industry trend makes it worse: moving from 800G to 1.6T and beyond raises the channel count per monitor, so the per-channel revisit interval lengthens precisely as the detuning tolerance tightens.

The observation this note exploits is that drift is **not independent across channels**. It decomposes into a common substrate/package mode, shared by every ring, and a local per-ring mode. The common mode is observable from *any* channel: an estimator that knows this structure refreshes it once per dwell period rather than once every 16.

3 · MODEL PARAMETERS (ALL PUBLIC)

Parameter	Value	Provenance
Effective thermo-optic coeff., $\mathrm{d}n_{\mathrm{eff}}/\mathrm{d}T$	$1.8\cdot10^{-4}$ K^{-1}	Material constant (Si, 220 nm strip with SiO_2 cladding)
Group index, n_g	4.00	220 nm Si strip, standard value
Resonance sensitivity, $\mathrm{d}\lambda/\mathrm{d}T$	69.8 pm/K	Derived: $\lambda\,(\mathrm{d}n_{\mathrm{eff}}/\mathrm{d}T)/n_g$ — not a free parameter
Loaded quality factor, Q	10000	Typical microring modulator/filter
Resonance linewidth, FWHM	155 pm	Derived: λ/Q
WDM channels in the bank	16	Representative of an 800G/1.6T link
Common mode (substrate/package)	$\sigma = 0.25$ K, $\tau = 20$ min	Package thermal constant; residual of the published 5–8 K inter-channel gradient
Local mode (per ring)	$\sigma = 0.10$ K, $\tau = 5$ min	Residual after the local heater loop ($\sim$1.5 % of the gradient)
Ambient/mounting drift	0.15 K/$\sqrt{\mathrm{h}}$	Slow random walk
Creep (relaxation/ageing)	0.05 K/h	Linear drift: 27.9 pm over 8 h
Resonance monitor noise	3.0 pm	2 % of FWHM — a conservative assumption, it favours the baseline

Table 1: 8 h horizon · 5 s actuation step · 60 Monte Carlo realisations × 16 rings · seed 42.

4 · STRATEGIES COMPARED

All five strategies receive exactly the same measurement budget and the same monitor noise. They differ only in what they do *between* measurements. C2 isolates the effect of predicting; C3 adds the effect of knowing the structure of the bank. C4 repeats C3 with a deliberately mis-specified model (thermal constants wrong by a factor of two, amplitudes mis-scaled, creep not modelled) as a robustness check.

Control strategy	RMS [pm]	Penalty [dB]	vs C1
C0 Static calibration (never updated)	35.1	0.809	0.47×
C1 Per-ring sample-and-hold ← standard	16.4	0.191	—
C2 Per-ring prediction (SISO)	15.5	0.170	1.06×
C4 Structured, mismatched model	10.8	0.084	1.52×
C3 Structured estimation of the bank	**9.1**	**0.059**	**1.80×**

Table 2: Nominal point (60 s dwell, 16 min revisit). The penalty is the drop-port penalty of a Lorentzian filter, $10\log_{10}(1+(2\delta/\mathrm{FWHM})^2)$; at half linewidth it is 3.01 dB.

5 · RESULTS

Prediction alone barely helps: the per-ring filter (C2) improves on sample-and-hold by a bare **6 %**, because in an isolated ring the residual error is dominated by unpredictable thermal innovation. The gain appears on exploiting the structure: joint estimation of the bank (C3) cuts RMS detuning from **16.4 pm to 9.1 pm**, a factor of 1.80×, of which 1.70× is attributable purely to cross-channel information. With the mismatched model (C4), 1.52× survives: the result does not depend on getting the parameters right.

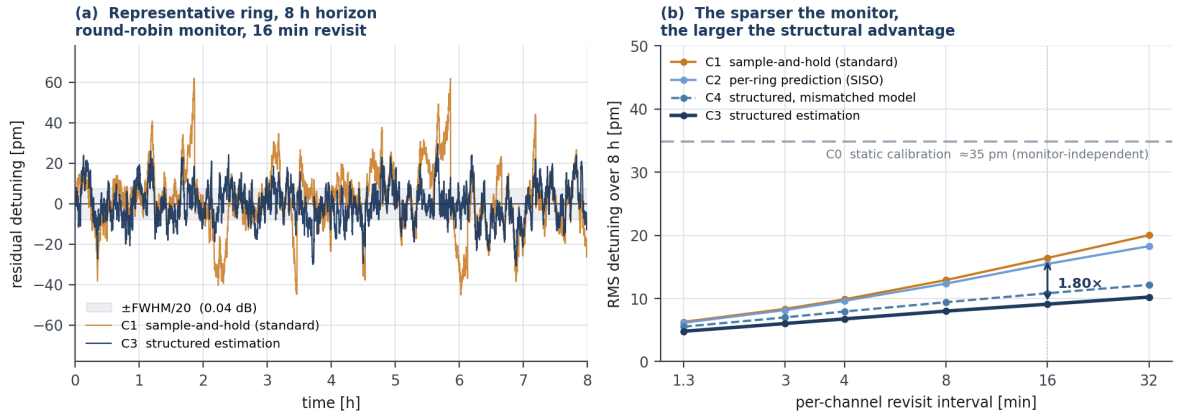


Figure 1: (a) Residual detuning of a representative ring over 8 h. (b) RMS against the per-channel revisit interval.

Panel (b) carries the result with the most industrial consequence: **the advantage grows as the monitor gets sparser**, from 1.31× at a 1.3 min revisit to 1.96× at 32 min. Read as a hardware equivalence, sample-and-hold would have to drop its dwell to 10 s —**6× more monitoring bandwidth**— to match what structured estimation achieves with the instrumentation already installed.

6 · WHAT THIS NOTE SHOWS, AND WHAT IT DOES NOT

- **It shows** that, under a physically reasonable drift model, there is margin recoverable in pure software at the control layer — no new chip, no extra heaters, no more cooling.
- **It shows** that the margin comes from the spatial structure of the problem rather than from the sophistication of the filter: it is the thermal coupling between channels that is being wasted.
- **It does not show** any result measured on silicon. This is simulation; it does not substitute for a characterisation campaign.
- **It does not predict** the improvement a given platform would obtain. The real factor depends on three quantities that can only come from measured characterisation data: the split between common and local mode, the effective thermal constants, and the actual monitor cadence. The values used here are plausible, not measured.
- **It simplifies** the physics: the common/local decomposition stands in for the full thermal coupling matrix of the bank. A real, characterised matrix admits more gain than shown, not less.

7 · NEXT STEP

The question this note leaves open —*where your platform sits on the curve in Figure 1(b)*— is answered with measured characterisation data under NDA. That is the object of the **Photonic Stability Diagnostic**: fitting this model to your platform's measured thermal constants and common/local split, and delivering the quantified improvement together with a working prototype of the compensation module.

If you would rather verify before contracting anything: supply an existing, anonymised drift record and I return a **sealed forecast** — committed by its SHA-256 hash before I see the second half of the record — at no cost and no obligation. The diagnostic is complete when a pre-agreed acceptance criterion is met; if it is not met, it is not invoiced.

One warning worth having, in case you would rather try it yourselves. The direct version of the method — one filter per ring, each fed only by its own measurements — returns **6 %**, and gives the impression that the approach does not work. The gain only appears once the bank is treated as a single jointly estimated system, and that is where the engineering lives.

The code reproducing every number and every figure in this note, together with the sealed-forecast tooling, is available as an open package under the MIT licence.

Kerym Makraini · Independent researcher, quantum photonics · Málaga, Spain
[email@domain.com] · [LinkedIn / web]