

Order-Preserving Channel Calibration of Kerr Comb-Driven Microresonator-Based DWDM Link

Yuyang Wang,^{1,*} Songli Wang,¹ Swarnava Sanyal,² Nathaniel Nauman,¹
Robert Parsons,¹ James Robinson,¹ Maarten Hattink,^{1,3} Kaylx Jang,¹
Asher Novick,^{1,3} Karl J. McNulty,¹ Xiang Meng,¹ Michal Lipson,^{1,2}
Alexander L. Gaeta,^{1,2} and Keren Bergman¹

¹ Department of Electrical Engineering, Columbia University, New York, NY 10027, USA

² Department of Applied Physics and Applied Mathematics, Columbia University, New York, NY 10027, USA

³ Now with Xscape Photonics Inc., Fort Lee, NJ 07024, USA

*yw3831@columbia.edu

Abstract: We experimentally validate a robust channel calibration algorithm for Kerr comb-driven microresonator-based DWDM links, which preserves the post-tuning resonator spectral order in the presence of resonance aliases within a comb spectrum spanning multiple resonator FSRs. © 2025 The Author(s)

1. Introduction

Silicon photonics microresonator-based dense wavelength-division multiplexing (DWDM) offers a cost-effective approach to massive wavelength parallelism, leveraging recent advances in normal-group velocity dispersion (GVD) Kerr frequency comb sources that can generate up to hundreds of precisely spaced channels from a single continuous-wave (CW) laser [1]. A scalable DWDM link architecture has been proposed to utilize the ultra-broad optical spectrum of the Kerr comb, typically spanning multiple resonator free spectral ranges (FSRs)—the spacing between two successive resonances of the same resonator [2]. Unlike conventional designs that confine the carrier wavelengths within a single resonator FSR, this *multi-FSR* architecture (Fig. 1a) allows for a manifold increase in usable optical bandwidth by strategically placing the unwanted resonances, i.e., *resonance aliases*, between comb channels while ensuring minimal crosstalk. As DWDM links feasibly scale beyond 64 channels per fiber [3], aligning resonator resonances with comb channels becomes increasingly challenging, particularly in the presence of in-band resonance aliases and wafer-scale process variations.

In this work, we address the challenge of robustly calibrating resonator resonances to comb channels in the multi-FSR regime, a crucial step for DWDM link *initialization*. Compared to wavelength locking—typically performed locally to each resonator through dedicated feedback loops [4]—initialization is a prerequisite and requires a global approach. This includes resolving the channel conflicts between multiple resonators and ensuring the transmitter (Tx) and receiver (Rx) resonators having the same spectral order. Although a shuffled Tx–Rx channel mapping can be corrected by an extra bit–re-shuffling component at the Rx electrical backend [5], it implies large area and latency overheads for high channel counts and thus is not preferred. Our proposed algorithm exploits a unique spectral property of the Kerr comb—*independent of any wavelength-specific information*—to achieve robust in-order channel calibration under wafer-scale process variations, with only a modest tuning range needed per resonator. This approach is thus pivotal for enabling scalable DWDM links in future computing systems with embedded photonics.

2. Channel Calibration Algorithm

Assuming a non-shuffled Tx–Rx channel mapping, transceiver channel calibration schemes vary in the level of order restrictions imposed on the Tx (or Rx) resonators post-tuning. The most stringent is *in-order* calibration, which strictly preserves the spectral order of the resonators, while the most relaxed is *shuffled* calibration, allowing for arbitrary channel re-mapping. In-order calibration is preferable for enabling wavelength-dependent optimizations such as channel equalization [6], but it is traditionally challenging due to the large tuning range required to compensate for the worst-case process variations. As a result, most single-FSR solutions, such as [5, 7], adopt an intermediate form called *cyclic* calibration—permitting a cyclic permutation of the spectral order by using resonance aliases for modulation/filtering—aiming at achieving a semi-deterministic channel mapping with reduced tuning distance. In the multi-FSR regime, however, the resonance aliases are by design distributed across the entire comb spectrum (Fig. 1b), making cyclic calibration impractical and leading to a shuffled result when similar strategies are applied. A robust calibration algorithm thus requires globally resolving the ambiguity between main resonances and resonance aliases as they are thermally tuned across the comb spectrum.

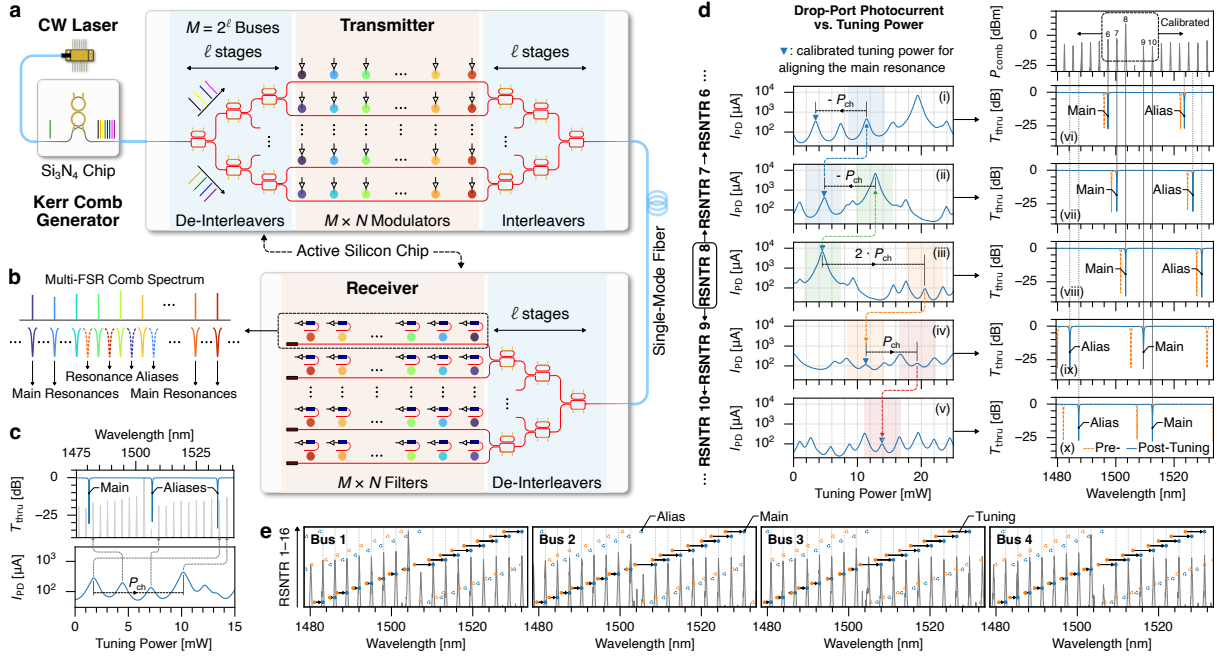


Fig. 1: (a) Scalable Kerr comb-driven DWDM architecture [2] featuring multiple buses of microresonators, using even-odd (de-)interleavers to expand the channel spacing on each bus. (b) The architecture allows for utilizing a comb spectrum much wider than the resonator FSR by placing resonance aliases between channels. (c) Thermal tuning red-shifts a resonator’s spectrum, producing a peak in the drop-port photocurrent trace when a resonance (or alias) captures a comb channel. Knowing the resonator tuning efficiency (nm/mW) and the comb channel spacing, P_{ch} can be calculated to identify two peaks corresponding to neighboring comb channels. (d) Simulated example of the *search* results and *match* process for a bus of 16 resonators. Resonator 8, whose main resonance ends up closest to the pump residue, is used as the anchor for sequentially calibrating the upstream and downstream resonators. (e) Simulated calibration result for a 4×16 link architecture, demonstrating fully in-order calibration for all 64 resonators and capability of skipping weak comb lines.

Our proposed calibration algorithm extends the two-phase approach in [8] to address the multi-FSR challenges. It consists of a *search* phase and a *match* phase. During the search phase, each resonator is thermally tuned across its tuning range, with upstream resonators (those physically closer to the light input) detuned. This scan should start from the maximum tuning power and move downward in light of the bistable states due to resonator self-heating [9]. The photocurrent from the drop-port photodetector (PD) of the resonator will exhibit multiple peaks (Fig. 1c), each indicating the alignment of either the main resonance or a resonance alias with a comb channel. The match phase then aims to identify the peak corresponding to the main resonance aligned with its target channel. Fig. 1d exemplifies this process with a modeled bus of 16 resonators. Here, we leverage the characteristic high-power line (pump residue) near the center of the normal-GVD Kerr comb spectrum. The resonator to align with the pump residue can be uniquely identified near the center of the bus by the highest PD reading with the least tuning power (Resonator 8 in Fig. 1d). Using this as an “anchor point”, we sequentially calibrate each upstream resonator by 1) matching the identified peak to a corresponding peak in the upstream resonator’s PD trace (Fig. 1d, (iii)→(ii) and (ii)→(i)), and 2) subtracting P_{ch} from the matched peak, where P_{ch} is the tuning power needed to shift the resonances by one comb channel spacing, as can be characterized from device measurement. The calibration of downstream resonators can be similarly achieved by first identifying the peak with a P_{ch} increment from the anchor point and then matching it with the downstream resonator’s PD trace (Fig. 1d, (iii)→(iv) and (iv)→(v)). The matching of the photocurrent peaks is automated by performing a moving-window cross-correlation between the anchor peak and the PD trace of interest. Note that we can also opt to skip a comb channel with a $2 \cdot P_{\text{ch}}$ increment/decrement, which is useful for avoiding known bad channels from the comb. Fig. 1e shows the calibration result for a simulated four-bus 4×16 link architecture, demonstrating fully in-order calibration for all 64 resonators while successfully skipping the weak lines of a measured 100 GHz comb. Consistent calibration results are obtained under resonance variations of a global 1σ of 0.65 nm (common across the bus) and a local 1σ of 0.17 nm (independent for each resonator), based on wafer-scale measurements. The algorithm requires only a fraction of the resonator FSR to be tuned, as the sequential calibration involves scanning only a few comb channels per resonator.

We also note the generality of the proposed algorithm for different link architectures and comb technologies: 1) For link architectures with a shared monitoring PD for the entire bus [8], the photocurrent peaks in Fig. 1c will appear as valleys, and our algorithm can be adjusted accordingly. 2) Despite leveraging the high-power pump residue for anchoring, the algorithm only requires identifying the highest PD reading of a near-center resonator to initiate the match phase, without needing precise information about the pump residue’s location or power. The algorithm is thus applicable to other comb technologies with deterministic spectral unevenness.

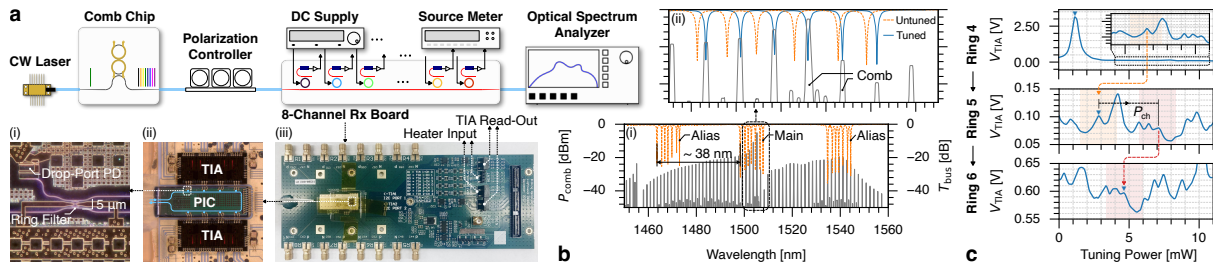


Fig. 2: (a) Experimental setup for validating the proposed calibration algorithm. (b) Comb and filter bus spectra showing multi-FSR challenges, and the calibration result for six filters. (c) TIA readings from the search phase showing matched peaks, underscoring the algorithm's practicality.

3. Experimental Validation

We experimentally validate our proposed channel calibration algorithm. The setup is shown in Fig. 2a, driven by a 200 GHz Kerr comb generated from silicon nitride dual-ring resonators pumped at 1,505 nm. The microresonator bus is emulated by an 8-channel, custom-designed, silicon microring-based Rx board (Fig. 2a(iii)). The board includes a photonic integrated circuit (PIC) with eight cascaded microring filters, optically packaged with an edge-coupled fiber array and 2.5D-integrated with commercial transimpedance amplifier (TIA) chips (Fig. 2a(ii)). Each microring filter has a drop-port PD for optical sensing (Fig. 2a(i)). The photocurrent from the PD is converted to a voltage signal, read by a source meter from a dedicated pin on the board. Thermal tuning of each filter is achieved by a direct current (DC) voltage, applied by a DC supply through the board. Both the DC supplies and the source meter are controlled by a computer which runs the calibration algorithm. An optical spectrum analyzer (OSA) is connected to the optical bus output for observing the calibration results, but is not required by the algorithm. Fig. 2b(i) plots the comb and the filter bus spectra before calibration, with notable resonance aliases within the comb band. Limited by equipment, we are able to control a maximum of six filters simultaneously. Fig. 2b(ii) shows the calibration results for the last six filters, achieving precise alignment to the comb channels in spectral order. In Fig. 2c, TIA readings from the search phase for the last three filters are plotted. The filters are measured with a tuning efficiency of 0.38 nm/mW, corresponding to a P_{ch} of 4 mW for 200 GHz comb spacing, as confirmed by the TIA traces. This cross-validation demonstrates the algorithm's practicality in real-world conditions.

4. Conclusion

We presented a robust channel calibration algorithm for Kerr comb-driven microresonator-based DWDM links, achieving fully in-order channel alignment in the multi-FSR regime. Simulation and experimental validation demonstrated robust calibration under wafer-scale process variations, with only a modest tuning range required. This approach thus provides a practical solution to the initialization of DWDM systems, paving the way for scalable photonic connectivity in future computing systems.

References

1. A. Rizzo *et al.*, "Massively scalable Kerr comb-driven silicon photonic link," *Nat. Photonics* **17**, 781–790 (2023).
2. A. Rizzo *et al.*, "Petabit-Scale Silicon Photonic Interconnects With Integrated Kerr Frequency Combs," *IEEE J. Sel. Top. Quantum Electron.* **29**, 1–20 (2023).
3. Y. Wang *et al.*, "Silicon Photonics Chip I/O for Ultra High-Bandwidth and Energy-Efficient Die-to-Die Connectivity," in *2024 IEEE Custom Integrated Circuits Conference (CICC)*, (IEEE, Denver, CO, USA, 2024), pp. 1–8.
4. K. Padmaraju *et al.*, "Wavelength Locking and Thermally Stabilizing Microring Resonators Using Dithering Signals," *J. Light. Technol.* **32**, 505–512 (2014).
5. M. Georgas *et al.*, "Addressing link-level design tradeoffs for integrated photonic interconnects," in *2011 IEEE Custom Integrated Circuits Conference (CICC)*, (IEEE, San Jose, CA, USA, 2011), pp. 1–8.
6. A. Novick *et al.*, "Integrated Photonic Resonant Modulator-Based Equalization and Optimization for DWDM," in *Optical Fiber Communication Conference (OFC) 2024*, (Optica Publishing Group, San Diego California, 2024), p. W1K.5.
7. A. V. Krishnamoorthy *et al.*, "Exploiting CMOS Manufacturing to Reduce Tuning Requirements for Resonant Optical Devices," *IEEE Photonics J.* **3**, 567–579 (2011).
8. M. Hattink *et al.*, "Automated tuning and channel selection for cascaded micro-ring resonators," in *Metro and Data Center Optical Networks and Short-Reach Links III*, (SPIE, San Francisco, United States, 2020), p. 24.
9. Q. Xu *et al.*, "Carrier-induced optical bistability in silicon ring resonators," *Opt. Lett.* **31**, 341–343 (2006).

Acknowledgements: This work was supported in part by the U.S. Defense Advanced Research Projects Agency (DARPA) under Photonics in the Package for Extreme Scalability (PIPES) program contract number HR00111920014 and in part by the Center for Ubiquitous Connectivity (CUBic) sponsored by the Semiconductor Research Corporation (SRC) and DARPA under the JUMP 2.0 program.