# Variation-Aware Adaptive Tuning for Nanophotonic Interconnects

Rui Wu[1,*], Chin-Hui Chen[2], Cheng Li[2], Tsung-Ching Huang[2], Fan Lan[3], Chong Zhang[1],
Yun Pan[4], John E. Bowers[1], Raymond G. Beausoleil[2], and Kwang-Ting Cheng[1]
[1]Department of Electrical & Computer Engineering, University of California, Santa Barbara, CA, US
[2]HP Labs, Hewlett-Packard Company, Palo Alto, CA, US
[3]College of Electrical Engineering, Zhejiang University, Hangzhou, Zhejiang, China
[4]Department of Information Science & Electronic Engineering, Zhejiang University, Hangzhou, Zhejiang, China
[*]ruiwu@ece.ucsb.edu

***Abstract*—Short-reach nanophotonic interconnects are promising to solve the communication bottleneck in data centers and chip-level scenarios. However, the nanophotonic interconnects are sensitive to process and thermal variations, especially for the microring structures, resulting in significant variation of an optical link's bit error rate (BER). In this paper, we propose a power-efficient adaptive tuning approach for nanophotonic interconnects to address the variation issues. During the adaptive tuning process, each nanophotonic interconnect is adaptively allocated *just enough* power to meet the BER requirement. The proposed adaptive tuning approach could reduce the photonic receiver power by 8% - 34% than the worst-case based fixed design while achieving the same BER. Our evaluation results show that the adaptive tuning approach scales well with the process variation, the thermal variation and the number of communication nodes, and can accommodate different types of NoC architectures and lasers.**

## I. Introduction

Nanophotonic interconnects provide high bandwidth, low energy consumption and low latency compared to traditional electrical interconnects. It becomes increasingly promising that the nanophotonic interconnects could replace the electrical links in short-reach applications, such as data centers, inter-chip, and intra-chip communications [1].

The microring resonator is widely used in many optical network-on-chip (NoC) architectures [2]–[7], because of its functional versatility, power efficiency, and compact footprint. However, the microring resonator is very sensitive to fabrication process variation and runtime thermal variation. As a result of the process and thermal variation effects, the resonance wavelength of the ring resonator deviates from the desirable carrier wavelength, which leads to performance degradation or even failure. This wavelength mismatch problem has been extensively studied: many power-efficient tuning and channel arrangement schemes have been proposed [8]–[10]; several feedback-based wavelength stabilization circuits have also been implemented [11], [12]. Existing work on thermal and process variations mainly focuses on the tuning of the resonance wavelength [13]–[17]. However, the process variation induced variations of quality factor (Q) and extinction ratio (ER) of the microring resonator have not been well studied.

We have fabricated batches of microring modulators and filters on 8 inch silicon-on-insulator (SOI) wafers at the CEA LETI foundry. The optical transmission spectra of the fabricated microring devices across the wafer are measured, from which the quality factor and the extinction ratio are extracted. We notice significant variations of Q and ER from both our fabricated microring resonators and literature [18], [19]. Meanwhile, the Q and ER are very important to determine the BER and power budget of an optical link [18]. Our simulation results show that the variations of Q and ER lead to significant variation of BER of the links. If the link design is targeted at the average performance of the devices, some of the links do not satisfy the BER requirement. Therefore, the variation effects of Q and ER must be carefully addressed.

A naive worst-case based fixed design can guarantee that most of the links satisfy the BER requirement. However, such a fixed design leads to excessive power consumption. In this paper, we propose a power-efficient adaptive tuning approach that tunes each link individually and allocates just enough power to meet the BER requirement. The adaptive tuning approach relies on on-chip fast BER estimation circuitry to monitor the link BER, and adaptively tunes either the laser or the photonic receiver to reach the target BER. We evaluate the power savings gained by the adaptive tuning approach with respect to different NoC architectures, variation values, and link configurations. Overall, the paper makes the following contributions:

- Characterizes the process variations of microring-based photonic devices using measured data.
- Demonstrates that the BER of the optical links could vary significantly due to the process variations of microring devices.
- Proposes an adaptive tuning approach that reduces the power consumption than the worst-case based fixed design with reasonable time and area overhead.
- Evaluates the adaptive tuning approach and demonstrates its scalability with respect to different levels of variations and various topologies.

## II. Background

Nanophotonic interconnects mainly consist of light sources, waveguides, photonic modulators and photonic receivers (Fig. 1). On-chip laser arrays and off-chip comb lasers are common choices for the light source [20]. In this work, we consider a distributed feedback (DFB) hybrid silicon laser as an example

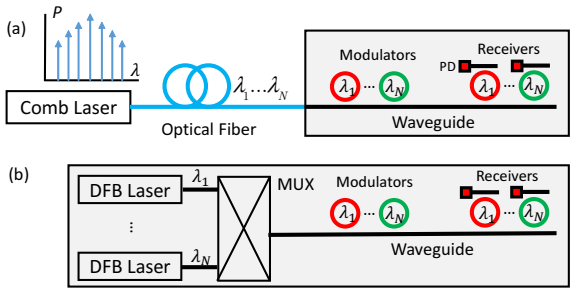


Fig. 1. Schematics of wavelength-division multiplexing (WDM) nanophotonic interconnects using (a) an off-chip comb laser or (b) an on-chip DFB laser array.

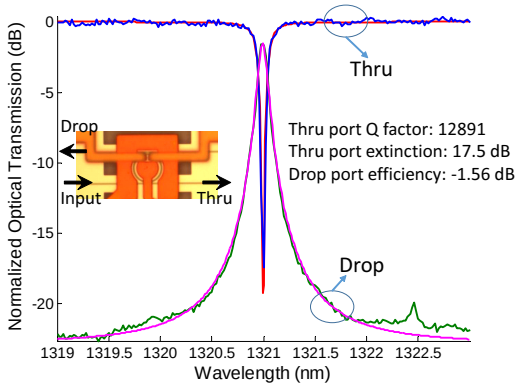


Fig. 2. The optical transmission of the through port and drop port of a microring (blue and green: measurement; red: model). The inset shows a microscopic image of a fabricated microring modulator.

of the on-chip single-wavelength laser [21], and a Gaussian shape comb laser for the off-chip laser [20]. Silicon waveguides are widely used to guide the light on SOI platforms. At the transmitter side, compact and energy-efficient microring modulators perform the on-off keying modulation of the light signal. At the receiver side, the light signal is redirected by the microring filter and sensed by the photodetector (PD).

The microring structure is critical in the nanophotonic interconnects, and a basic model is introduced here. When an integer number of the incident light wavelength fits the microring perimeter, the microring is on-resonance. At the on-resonance state, the through port power reaches its minima and the drop port power reaches its maxima. The optical transmission spectrum of the through port and the drop port can be described by the Lorentzian shape models [22] (Fig. 2):

$$T_{thru}(\lambda) = 1 - \frac{A_{thru}}{1 + (2Q \cdot (\lambda - \lambda_r)/\lambda_r)^2}$$
$$T_{drop}(\lambda) = \frac{A_{drop}}{1 + (2Q \cdot (\lambda - \lambda_r)/\lambda_r)^2} \quad (1)$$

where $\lambda_r$ is the the microring's resonance wavelength; $Q$ is the microring's quality factor; $A_{thru}$ is a parameter that is related to the microring's extinction ratio: $ER = 1/(1 - A_{thru})$. Furthermore, we denote the optical transmission at the on-(off-) resonance state as $T_{on} = T(\lambda = \lambda_r)$ ($T_{off} = T(\lambda = \lambda_r + \Delta\lambda)$), where $\Delta\lambda$ is the wavelength detuning for the off-resonance state. In this way, when the microring functions as a modulator, the optical transmission at logic "0" and logic "1" are $T_0 = T_{thru,on}$ and $T_1 = T_{thru,off}$, respectively. When the microring functions as a filter, the input port to drop port insertion loss (the drop port efficiency) is $T_{drop,on}$. Microring structures could also be used to build optical routers [23], [24]. For instance, in the five-port optical router reported in [23], the west-to-east insertion loss could be expressed as $T_{drop,on} \cdot T_{thru,off}^4 \cdot T_{WG\_loss}$.

Based on the theoretical device models above, we obtain the equation for the bit-error-rate (BER), a widely used figure of merit for communication quality:

$$\text{BER} = \frac{1}{2}\text{erfc}\left(\frac{z}{\sqrt{2}}\right), \quad z = P_{laser} \prod_i IL_i \cdot R_{pd} \frac{T_1 - T_0}{\sigma_1 + \sigma_0} \quad (2)$$

where $P_{laser}$ is the laser output power; $IL_i$ is the insertion loss of the photonic component $i$ along the optical path; $R_{pd}$ is the responsivity of the photodetector; $\sigma_1(\sigma_0)$ is the standard deviation of the logic "1" ("0") corresponding noise.

## III. Variation Challenges

Similar to deep submicron electronic devices, the nanophotonic devices (e.g., microring modulators, microring filters, grating couplers, photodetectors) also suffer from significant process variations [19]. In this work, we mainly focus on the severe variation effects in microring based devices, while our variation-aware analysis and design can also accommodate variations in other types of devices.

The microring structure is very sensitive to runtime thermal variation and fabrication-induced process variation. Due to the thermal variation effect, the optical transmission spectrum redshifts as the temperature rises. Due to the process variation effect, the device geometry and the waveguide sidewall roughness vary in the fabrication process. Consequently, the $\lambda_r$, Q and ER deviate from the designed values. Both thermal and process variation effects will cause the mismatch between the resonance wavelength and the carrier wavelength. Many tuning schemes and circuits have been proposed to address the wavelength mismatch problem [8], [9], [11], [12], [25]. However, few tuning schemes take into account the variations of the Q and the ER.

Here we characterize the variations of the Q and the ER based on measured results. Fig. 3 plots our wafer-scale inter-die measured data of the fabricated microring devices, together with the intra-die variation testing result reported in [18]. Both the inter-die and intra-die measured results show wide distribution ranges of Q and ER, which may have great impact on the communication BER. From the histograms, one can see that the distribution of the parameter $A$ and the Q approximately follow normal distributions.

The electrical tuning is usually utilized to compensate for the wavelength mismatch because it's more power efficient than the thermal tuning [13], [26]. However, the Q and the ER degrade significantly when the tuning voltage is applied to the microring resonator (Fig. 4). The severe degradation of Q and ER caused by the electrical tuning may greatly deteriorate the communication BER, which also needs to be carefully considered during the link analysis.

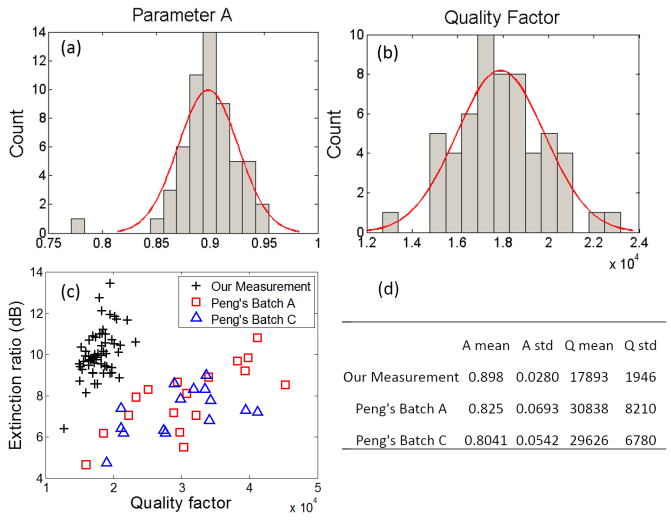


| | A mean | A std | Q mean | Q std |
|---|---|---|---|---|
| Our Measurement | 0.898 | 0.0280 | 17893 | 1946 |
| Peng's Batch A | 0.825 | 0.0693 | 30838 | 8210 |
| Peng's Batch C | 0.8041 | 0.0542 | 29626 | 6780 |

Fig. 3. (a)(b) Histograms of parameters A and Q of our inter-die measurement results; (c) Scatter plot of our inter-die measurement and Peng's intra-die measurement in [18]; (d) Mean and standard deviation (std) of the three measured data sets.

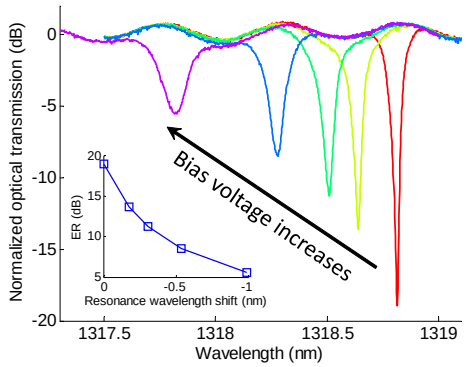


Fig. 4. The measured spectra series of a microring modulator at different bias voltages.

The Q and the ER are important to determine the microring's optical transmission and the BER (Eq. 1 and 2). The large variations of Q and ER may result in significant variation of the BER. We perform Monte Carlo simulations of a simple single-writer single-reader (SWSR) link [2]. The process variation statistics of our fabricated devices are used (the first line in Fig. 3d), which has better ER and uniformity than the Peng's devices in [18]. The simulation results in Fig. 5 show that the BER has a wide distribution range. About half of the links do not satisfy the BER requirement if the link design is based on the mean parameter values of the devices. Naively, a worst-case based fixed design could guarantee that most of the links satisfy the BER requirement (the red line in Fig. 5). However, such a fixed design requires excessive power consumption. For instance, in our simulation the laser output power needs to be increased by 20% to guarantee that 99% of the fabricated links satisfy the $10^{-12}$ BER requirement.

## IV. Adaptive Tuning Approach

Instead of the power-consuming, worst-cased based fixed design, we propose an adaptive tuning approach that tunes

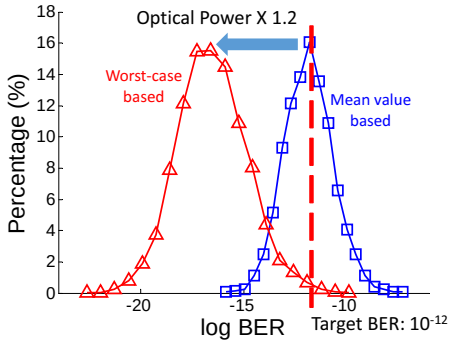


Fig. 5. The BER distribution in the presence of process variations and the electrical tuning. The blue line represents the link configuration based on the means of the device parameters. The red line enhances the link optical power by 1.2X

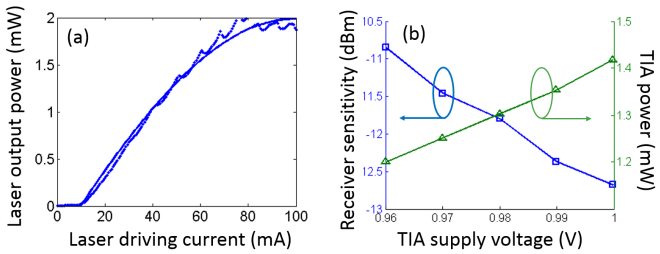


Fig. 6. (a) The laser output power as a function of the driving current of a DFB laser [21] (dots: measurement, line: model); (b) The sensitivity and power consumption versus supply voltage of an adaptive photonic receiver (data from [11]).

each optical link individually to meet the BER requirement. At a given BER and data rate, the minimum required optical modulation amplitude (OMA) is determined by the receiver sensitivity $P_{sense}$:

$$P_{laser} \prod_i IL_i \cdot (T_1 - T_0) = P_{sense} \tag{3}$$

If the $IL_i$, $T_1$, and $T_0$ vary due to process and thermal variations, we could tune either the laser output power $P_{laser}$ or the receiver sensitivity $P_{sense}$ to satisfy the above equation. Intuitively, we could tune the DFB laser's output power by varying its driving current (Fig. 6 a). However, for interconnect schemes using an off-chip comb lasers, it's inefficient to tune the laser's output power. This is because a comb laser has a fixed optical spectrum distribution; and individual wavelength cannot be tuned independently.

Fortunately, we notice an effective mechanism to trade off the power consumption for the receiver sensitivity (Fig. 6 b). The supply voltage of the receiver circuitry, i.e. the trans-impedance amplifier (TIA), has a significant impact on the gain, bandwidth, and noise performance [11]. As the TIA supply voltage increases, the circuitry consumes more power and the receiver achieves a better sensitivity (Fig. 6 b). Another benefit of tuning the receiver is allowing the sharing of on-chip lasers [27]. In summary, for optical links using on-chip DFB lasers, we could tune either the laser output power or the receiver sensitivity. For links using off-chip comb lasers, we could only tune the receiver sensitivity.

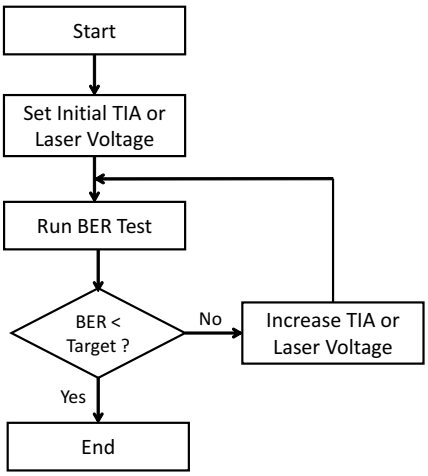


Fig. 7. The adaptive tuning flow.

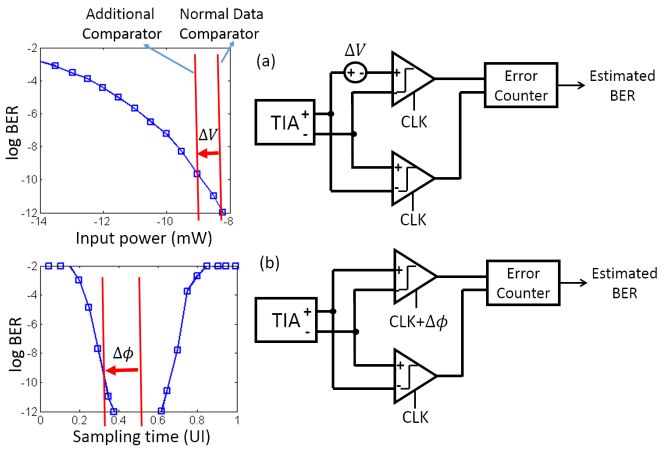


Fig. 8. The two fast BER estimation methods: (a) offset the voltage; (b) offset the sampling time. The data in the left two figures are from [30].

Fig. 7 illustrates the flow of the adaptive tuning approach. At the beginning, a specific writer-reader communication pair is activated. The TIA supply voltage or the laser driving voltage is set to a relatively low value based on the best-case device parameters. Then an on-chip BER testing circuitry performs the BER test. The TIA or laser voltage is gradually increased until the BER is below a pre-set target (eg. $10^{-12}$). Finally, the *just enough* TIA or laser driving voltage is stored in the flash memory as a lookup table (LUT). At runtime, the laser or the TIA voltage is configured based on the corresponding data stored in the LUT. This adaptive tuning flow could be activated before shipment, after deployment, or whenever an abnormal error rate is observed by higher level blocks (e.g., by using parity check) during operation . It should be noted that when the electrical tuning is applied, the TIA or laser power should also be increased accordingly at runtime to compensate for the degradation of Q and ER.

One of the key enabler in the adaptive tuning scheme is the fast on-chip BER testing circuitry. The brute-force BER testing method is unaffordably time-consuming (over 100 seconds for $10^{-12}$ BER at 10Gbps data rate) for the proposed adaptive tuning. Fortunately, there are fast BER estimation methods that leverage voltage offsetting or sampling time offsetting [28], [29]. As illustrated in Fig. 8, the BER decreases when the received power decreases or when the sampling time deviates from the ideal sampling point. Therefore, the voltage offsetting method or the sampling time offsetting method could be leveraged to accelerate the BER test. These two methods use an additional comparator that intentionally decreases the received power or deviates the sampling time. By comparing the additional comparator's output with that of the normal data comparator, eye closure could be detected much faster. For instance, if our target BER is $10^{-12}$, the BER after voltage or sampling time offsetting can be intentionally increased to about $10^{-10}$. It takes only several seconds to estimate such an increased BER for a 10 Gbps link.

There are two overheads associated with the adaptive tuning approach: the hardware overhead and the tuning time overhead. The hardware overhead, mainly in the BER monitor, includes an additional comparator with offset control and a small logic circuit, which takes about 30 $\mu m$ x 30 $\mu m$ area for a 65 nm CMOS technology [28]. For a multi-receiver link, its hardware cost can be amortized by sharing the BER monitor, as the multiple receivers on one link cannot receive signal simultaneously. In this way, for a 64-cluster crossbar with 64 WDM channels, the total area cost of the BER monitors is only 1.0% of the chip area (366.1 $mm^2$ in [27]). The tuning power overhead is avoided at runtime as the BER monitor circuit is switched off after the tuning process is complete. The tuning time overhead for a link is proportional to the number of communication pairs times the BER estimation time. For a WDM link, all wavelength channels could be tuned concurrently. For instance, assuming the BER estimation time is about 5 s, the tuning time for a many-writer single-reader (MWSR) link with 64 clusters (e.g., Corona [3]) is 320 s, which is reasonable for a one-time overhead. Overall, the hardware and time overheads are reasonable for practical applications.

## V. Evaluations

In this section, we perform simulations and analysis of several common photonic NoC architectures to evaluate the power saving gained by the proposed adaptive tuning approach. Several representative types of link structures are identified (as illustrated in Fig. 9) for experiments among the common photonic NoC architectures:

- Single-writer single-reader (**SWSR**): The SWSR point-to-pint link is used in the three-stage Clos network [2]. The Clos network uses SWSR optical links for stage-to-stage communication and electrical routers for routing.
- Many-writer single-reader (**MWSR**): The MWSR optical links can be used to construct optical crossbars for optical NoC. For instance, the Corona architecture replicates the MWSR channel 64 times to fully connect the 64 clusters [3]. For the adaptive tuning, each possible communication pair is individually tuned and the just enough TIA supply voltages are stored in a LUT.
- Single-writer many-reader (**SWMR**): Similar to the MWSR structure, the SWMR links can also be used

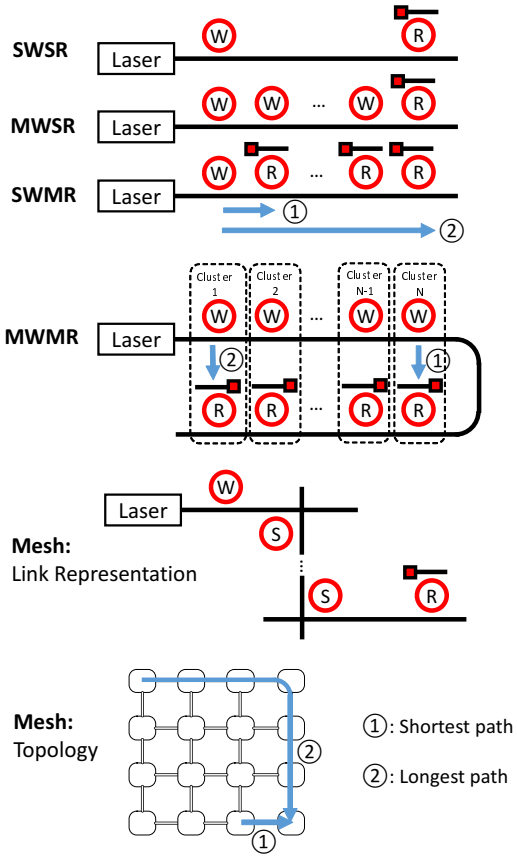


Fig. 9. Five types of link structures. Notions: W: writer or photonic modulator; R: reader or photonic receiver; S: photonic switch or router. In a WDM system, the laser represents a multi-wavelength light source; each microring represents a microring array for multiple wavelength channels

in crossbar structures. For instance, the Firefly architecture uses a SWMR-based crossbar for inter-cluster communication and a concentrated mesh for intra-cluster [5]. Additionally, the SWMR structure is used in an optical bus-based NoC architecture to broadcast the optical signal [4].

- Many-writer many-reader (**MWMR**): LumiNOC leverages the MWMR structure for subnet design, where each cluster is connected with a writer and a reader such that any cluster can communicate with any other cluster [6].
- Optical router based **mesh**: Microring-based optical routers are used in mesh-based NoC architectures in a manner of circuit switching. For instance, Petracca et al. proposed a non-blocking mesh NoC architecture using 4x4 optical routers [7].

### A. Experimental Setup

We perform Monte Carlo simulations to calculate the power consumptions by the fixed design and by the adaptively tunable design. The process variation statistics of parameters $A$ and Q in the first line of Fig. 3 d are used to generate random instances of microring devices. The standard deviation of resonance wavelength caused by the process variation is 0.44 nm [13]. The channel spacing is set to 0.9 nm (160 GHz) to match the grid of the laser [31]. Since the electrical tuning is

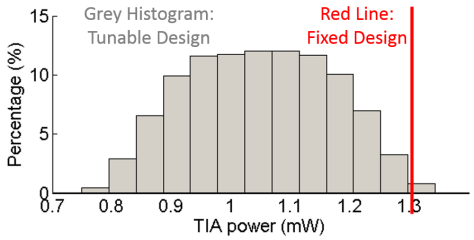


Fig. 10. The simulated TIA power consumption for the fixed design and the tunable design.

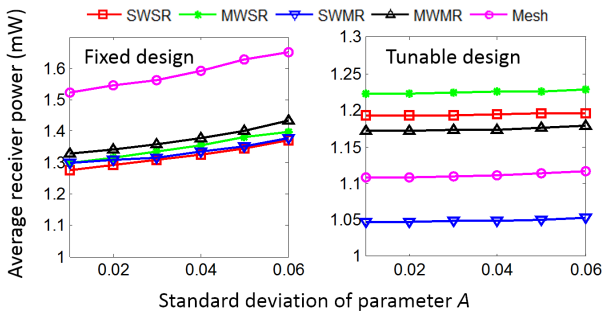


Fig. 11. The average receiver power consumption at different process variations and for different link structures.

power-efficient and the electrical tuning range (1 nm in Fig. 4) can cover the channel spacing (0.9 nm), we therefore adopt the electrical tuning with the channel remapping (or reshuffling) technique to compensate for the wavelength mismatch [8], [9]. Other than specified for parameter sweeps, the number of clusters is set to 16; the maximum temptation variation is set to 17°C [13]; the laser type is on-chip DFB laser. The length of the longest communication path is set to 4 cm to accommodate the chip area (366.1 $mm^2$ in [27]). The waveguide loss is assumed to be 0.74 dB/cm [32].

The fixed design uses the same TIA supply voltage for all photonic receivers; the yield target is set to 99%. The adaptive tuning leverages the sensitivity-adaptive photonic receiver illustrated in Fig. 6 b, and configures each TIA's supply voltage individually depending on the present communication pair. Each writer (or reader) in the many-writer (or reader) structure is assumed to have the same probability to write (or read). The average power consumption of a TIA for the fixed design and the tunable design are reported in the simulations. Fig. 10 shows the Monte Carlo simulation result of a SWMR link, where the fixed design has to set the TIA power as high as 1.32 mW while the adaptively tunable design achieves an average TIA power of 1.05 mW.

### B. Process and Thermal Variations

From Fig. 3 (c)(d), one can see that the process variation statistics are very distinct for different fabrication processes. In order to evaluate the performance of the adaptive tuning approach at different process variation levels, we sweep the standard deviation of the parameter $A$, which is an important factor to determine the link's BER. Fig. 11 shows that the average receiver power by the adaptive tuning is decoupled from the process variation.

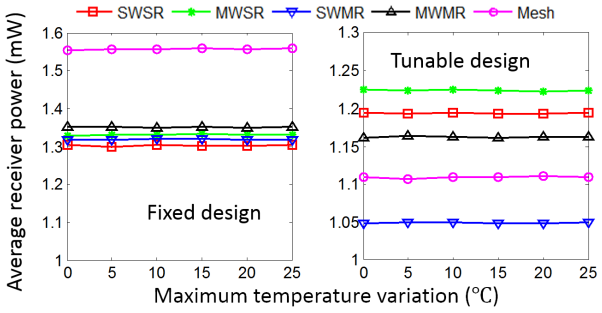


Fig. 12. The average receiver power consumption at different temperature variations and for different link structures.

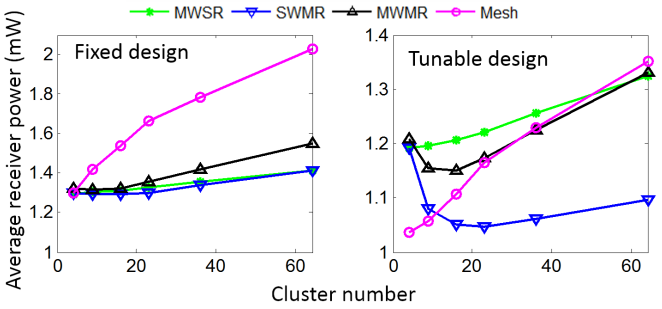


Fig. 13. The average receiver power consumption at different numbers of clusters and for different link structures.

Temperature variations of an NoC is design- and workload-dependent. We simulate the average receiver power at different maximum temperature variation values. The simulation results in Fig. 12 show that the receiver power of both the fixed design and the tunable design are decoupled from the temperature variation. This is because the channel remapping technique [26] decouples the required tuning distance from the the thermal variation when the process variation ($3\sigma = 1.32nm$) is greater than the channel spacing (0.9 nm). This decoupled phenomenon has also been observed in [9], [26].

### C. Number of Clusters

We also evaluate the power savings with respect to different numbers of clusters. Fig. 13 shows that, for SWMR, MWMR, and Mesh links, the power saving will be greater for a design with more clusters. This is due to the fact that the communication path lengths are highly non-uniform in SWMR, MWMR, and Mesh links, as illustrated by the shortest and longest paths in Fig. 9. The shortest and longest paths potentially lead to very distinct best-case and worst-case link performance [33]. In other words, the non-uniformity of the communication path lengths, together with the process and thermal variations, contribute to the power savings.

### D. Type of Lasers

Previous simulations all assume the links use on-chip DFB lasers as the light source. Comb lasers intrinsically have non-uniform output power at different wavelengths, for which the adaptive tuning should achieve even greater power saving. We simulate the comb laser based optical links with 16 channels. The output spectrum of a Gaussian comb laser is described in

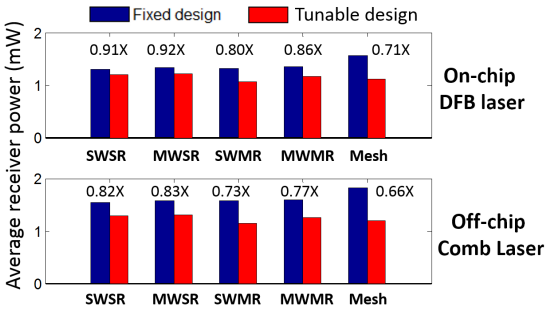


Fig. 14. The average receiver power consumption for different laser types and for different link structures. The numbers denote the ratios of the average receiver power by the tunable design to the fixed design.

[20]. In the simulations, we assume that the output power of the comb laser at its center wavelength (the maxima) equals the DFB laser output used in previous simulations plus the fiber-to-chip coupling loss. The simulation results confirm that the adaptive tuning leads to greater power savings for comb lasers based links (Fig. 14).

In summary, the adaptive tuning approach scales well with the process variation, the thermal variation, and the number of clusters, especially for comb lasers based links and for link structures with non-uniform path lengths.

## VI. Conclusion

Microring resonator based nanophotonic interconnects are very sensitive to process and thermal variations. In this paper, we model the microring based photonic devices and their variation effects using the measured data. Taking into account the process and thermal variation effects, our simulations show that the BER of optical links has significant variation. Since the worst-case based fixed design consumes much excessive power, we propose a novel power-efficient adaptive tuning approach. The proposed approach could tune each link individually and allocate just enough power to meet the BER requirement. This approach offers good power efficiency with reasonable area and time overhead. Our simulation and analysis demonstrate that the proposed adaptive tuning approach scales well with respect to different process variations, thermal variations, numbers of clusters, and laser types. Particularly, the adaptive tuning could save more power for link structures with non-uniform communication path lengths and/or using comb lasers.

## Acknowledgment

R. Wu, J. E. Bowers, and K.-T. Cheng acknowledge the Semiconductor Research Corporation (SRC) for the support.

## References

[1] R. G. Beausoleil, M. McLaren, and N. P. Jouppi, "Photonic architectures for high-performance data centers," *Selected Topics in Quantum Electronics, IEEE Journal of*, vol. 19, no. 2, pp. 3700109–3700109, 2013.

[2] A. Joshi, C. Batten, Y.-J. Kwon, S. Beamer, I. Shamim, K. Asanovic, and V. Stojanovic, "Silicon-photonic clos networks for global on-chip communication," in *International Symposium on Networks-on-Chip*, 2009.

[3] D. Vantrease, R. Schreiber, M. Monchiero, M. McLaren, N. P. Jouppi, M. Fiorentino, A. Davis, N. Binkert, R. G. Beausoleil, and J. H. Ahn, "Corona: System implications of emerging nanophotonic technology," in *Computer Architecture, International Symposium on*, 2008.

[4] N. Kirman, M. Kirman, R. K. Dokania, J. F. Martinez, A. B. Apsel, M. A. Watkins, and D. H. Albonesi, "Leveraging optical technology in future bus-based chip multiprocessors," in *International Symposium on Microarchitecture*, 2006.

[5] Y. Pan, P. Kumar, J. Kim, G. Memik, Y. Zhang, and A. Choudhary, "Firefly: illuminating future network-on-chip with nanophotonics," in *ACM SIGARCH Computer Architecture News*, 2009.

[6] C. Li, M. Browning, P. V. Gratz, and S. Palermo, "Luminoc: A power-efficient, high-performance, photonic network-on-chip," *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on*, vol. 33, no. 6, pp. 826–838, 2014.

[7] M. Petracca, B. G. Lee, K. Bergman, and L. P. Carloni, "Design exploration of optical interconnection networks for chip multiprocessors," in *High Performance Interconnects (HOTI), IEEE Symposium on*, 2008.

[8] Y. Zheng, P. Lisherness, M. Gao, J. Bovington, K.-T. Cheng, H. Wang, and S. Yang, "Power-efficient calibration and reconfiguration for optical network-on-chip," *Journal of Optical Communications and Networking*, vol. 4, no. 12, pp. 955–966, 2012.

[9] M. Georgas, J. Leu, B. Moss, C. Sun, and V. Stojanovic, "Addressing link-level design tradeoffs for integrated photonic interconnects," in *Custom Integrated Circuits Conference (CICC)*, pp. 1–8, 2011.

[10] Y. Xu, J. Yang, and R. Melhem, "Tolerating process variations in nanophotonic on-chip networks," in *International Symposium on Computer Architecture*, 2012.

[11] C. Li, R. Bai, A. Shafik, E. Z. Tabasy, G. Tang, C. Ma, C.-H. Chen, Z. Peng, M. Fiorentino, P. Chiang, *et al.*, "A ring-resonator-based silicon photonics transceiver with bias-based wavelength stabilization and adaptive-power-sensitivity receiver," in *International Solid-State Circuits Conference (ISSCC)*, 2013.

[12] X. Zheng, E. Chang, P. Amberg, I. Shubin, J. Lexau, F. Liu, H. Thacker, S. S. Djordjevic, S. Lin, Y. Luo, *et al.*, "A high-speed, tunable silicon photonic ring modulator integrated with ultra-efficient active wavelength control," *Optics express*, vol. 22, no. 10, pp. 12628–12633, 2014.

[13] Z. Li, M. Mohamed, X. Chen, E. Dudley, K. Meng, L. Shang, A. R. Mickelson, R. Joseph, M. Vachharajani, B. Schwartz, *et al.*, "Reliability modeling and management of nanophotonic on-chip networks," *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, vol. 20, no. 1, pp. 98–111, 2012.

[14] Y. Ye, Z. Wang, P. Yang, J. Xu, X. Wu, X. Wang, M. Nikdast, and L. Duong, "System-level modeling and analysis of thermal effects in wdm-based optical networks-on-chip," *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on*, vol. 33, no. 11, pp. 1718–1731, 2014.

[15] D. Ding, B. Yu, and D. Z. Pan, "Glow: A global router for low-power thermal-reliable interconnect synthesis using photonic wavelength multiplexing," in *Design Automation Conference (ASP-DAC),Asia and South Pacific*, 2012.

[16] J. Chan, G. Hendry, A. Biberman, K. Bergman, and L. P. Carloni, "PhoenixSim: A simulator for physical-layer analysis of chip-scale photonic interconnection networks," in *Design, Automation & Test in Europe Conference & Exhibition (DATE)*, 2010.

[17] C. Condrat, P. Kalla, and S. Blair, "Thermal-aware synthesis of integrated photonic ring resonators," in *International Conference on Computer-Aided Design*, 2014.

[18] Z. Peng, D. Fattal, M. Fiorentino, and R. Beausoleil, "Fabrication variations in soi microrings for dwdm networks," in *Group IV Photonics (GFP)*, 2010.

[19] P. P. Absil, P. Verheyen, P. De Heyn, M. Pantouvaki, G. Lepage, J. De Coster, and J. Van Campenhout, "Silicon photonics integrated circuits: a manufacturing platform for high density, low power optical i/os," *Optics Express*, vol. 23, no. 7, pp. 9369–9378, 2015.

[20] M. J. Heck and J. E. Bowers, "Energy efficient and energy proportional optical interconnects for multi-core processors: Driving the need for on-chip sources," *Selected Topics in Quantum Electronics, IEEE Journal of*, vol. 20, no. 4, pp. 332–343, 2014.

[21] C. Zhang, S. Srinivasan, Y. Tang, M. J. Heck, M. L. Davenport, and J. E. Bowers, "Low threshold and high speed short cavity distributed feedback hybrid silicon lasers," *Optics express*, vol. 22, no. 9, pp. 10202–10209, 2014.

[22] R. Wu, C.-H. Chen, J.-M. Fedeli, M. Fournier, K.-T. Cheng, and R. G. Beausoleil, "Compact models for carrier-injection silicon microring modulators," *Optics Express*, vol. 23, no. 12, pp. 15545–15554, 2015.

[23] R. Ji, L. Yang, L. Zhang, Y. Tian, J. Ding, H. Chen, Y. Lu, P. Zhou, and W. Zhu, "Five-port optical router for photonic networks-on-chip," *Optics express*, vol. 19, no. 21, pp. 20258–20268, 2011.

[24] N. Sherwood-Droz, H. Wang, L. Chen, B. G. Lee, A. Biberman, K. Bergman, and M. Lipson, "Optical 4x4 hitless slicon router for optical networks-on-chip (noc)," *Optics express*, vol. 16, no. 20, pp. 15915–15922, 2008.

[25] M. Valad Beigi and G. Memik, "Min: a power efficient mechanism to mitigate the impact of process variations on nanophotonic networks," in *International symposium on Low power electronics and design*, pp. 299–302, 2014.

[26] Y. Zheng, P. Lisherness, M. Gao, J. Bovington, S. Yang, and K.-T. Cheng, "Power-efficient calibration and reconfiguration for on-chip optical communication," in *Design, Automation & Test in Europe Conference & Exhibition (DATE)*, 2012.

[27] C. Chen, T. Zhang, P. Contu, J. Klamkin, A. K. Coskun, and A. Joshi, "Sharing and placement of on-chip laser sources in silicon-photonic nocs," in *Networks-on-Chip (NoCS)*, 2014.

[28] C. Li, R. Bai, A. Shafik, E. Tabasy, B. Wang, G. Tang, C. Ma, C.-H. Chen, Z. Peng, M. Fiorentino, R. Beausoleil, P. Chiang, and S. Palermo, "Silicon photonic transceiver circuits with microring resonator bias-based wavelength stabilization in 65 nm cmos," *Solid-State Circuits, IEEE Journal of*, vol. 49, pp. 1419–1436, June 2014.

[29] D. Hong and K.-T. Cheng, *Efficient test methodologies for high-speed serial links*. Springer Science & Business Media, 2009.

[30] K. Yu, C.-H. Chen, A. Titriku, A. Shafik, M. Fiorentino, P. Y. Chiang, S. Palermo, *et al.*, "25gb/s hybrid-integrated silicon photonic receiver with microring wavelength stabilization," in *Optical Fiber Communication Conference*, 2015.

[31] C.-H. Chen, C. Li, A. Shafik, M. Fiorentino, P. Chiang, S. Palermo, and R. Beausoleil, "A wdm silicon photonic transmitter based on carrier-injection microring modulators," in *Optical Interconnects Conference*, 2014.

[32] S. K. Selvaraja, P. De Heyn, G. Winroth, P. Ong, G. Lepage, C. Cailler, A. Rigny, K. Bourdelle, D. VanThourhout, J. Van Campenhout, *et al.*, "Highly uniform and low-loss passive silicon photonics devices using a 300mm cmos platform," in *Optical Fiber Communication Conference*, 2014.

[33] Y. Xie, M. Nikdast, J. Xu, X. Wu, W. Zhang, Y. Ye, X. Wang, Z. Wang, and W. Liu, "Formal worst-case analysis of crosstalk noise in mesh-based optical networks-on-chip," *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, vol. 21, no. 10, pp. 1823–1836, 2013.