

Power-Efficient Calibration and Reconfiguration for On-Chip Optical Communication

Yan Zheng^{1,2}, Peter Lisherness², Ming Gao², Jock Bovington², Shiyuan Yang¹, and Kwang-Ting Cheng²

1. Department of Automation from Tsinghua University, Beijing 100084, China,

2. Department of Electrical and Computer Engineering, University of California, Santa Barbara, CA 93106

e-mail : zheng-y03@mails.thu.edu.cn, {peter, mgao, jock}@ece.ucsb.edu, ysy-dau@tsinghua.edu.cn, timcheng@ece.ucsb.edu

Abstract—On-chip optical communication infrastructure has been proposed to provide higher bandwidth and lower power consumption for the next-generation high performance multi-core systems. Online calibration of these optical components is essential to building a robust optical communication system, which is highly sensitive to process and thermal variation. However, the power consumption of existing tuning methods to properly calibrate the optical devices would be prohibitively high. We propose two calibration and reconfiguration techniques that can significantly reduce the worst case tuning power of a ring-resonator-based optical modulator: 1) a channel re-mapping scheme, with sub-channel redundant resonators, which results in significant reduction in the amount of required tuning, typically within the capability of voltage based tuning, and 2) a dynamic feedback calibration mechanism used to compensate for both process and thermal variations of the resonators. Simulation results demonstrate that these techniques can achieve a 48X reduction in tuning power - less than 10W for a network with 1-million ring resonators.

I. INTRODUCTION

On-chip optical interconnects providing high bandwidth density and low power consumption have been proposed to meet future on-chip systems' bandwidth and performance requirements [1]. Being able to use CMOS-compatible fabrication processes also makes this technology more feasible in the future [2].

A basic building block of such systems is the micro-ring resonator [3], which can function as a modulator [4], filter [5], switch [6], or laser [7]. Due to its wavelength-specificity, Wavelength Division Multiplexing (WDM) can be used without the need for optical mux/demux, providing multiple independent and simultaneous data channels on a single waveguide and increasing bandwidth density significantly [8] [9]. One example of an optically connected design utilizing ring resonators is Corona [10], a 256-core 20 teraflop system constructed with 3D die stacking, targeting the 16nm process node in 2017 [11].

These ring resonators suffer from two kinds of variation: static fabrication process variation and dynamic thermal variation. WDM requires all the ring resonators working on the same physical channel to have the same resonant wavelength, which must also match the wavelength of the carrier light source for that channel. Variations will shift the resonant wavelength and cause a mismatch between the resonant wavelengths of the rings and the carrier, leading to performance degradation, or even complete failure [12].

There are two tuning methods to compensate for these variations: bias voltage tuning [13] and thermal tuning with local heaters [14]. The power consumed by tuning is highly

dependent on the magnitude of the deviation in resonant wavelength from the carrier wavelength. We simulate the average difference, or "tuning distance", in this paper, and show that a naive direct tuning scheme will consume excessive power, compromising the practicality of on-chip optical communication.

In this paper, we propose an online reconfiguration and calibration technique for on-chip optical communications to limit the average tuning distance and reduce tuning power consumption.

The contributions of this paper:

1. We propose a channel-remapping scheme with sub-channel redundant resonators that opportunistically reassigns logical channels to the carriers through one of the redundant rings to decrease the worst case tuning distance from one FSR (Free Spectral Range) to less than one channel so that the calibration can be achieved solely by the bias voltage based tuning. Our simulation results show that we can decrease the worst-case tuning distance of one ring to as little as 2.1% of the original, leading to a similar reduction in tuning power.

2. We also propose a self-calibration mechanism to compensate for the effects of process and thermal variation on resonant wavelength with affordable cost. This circuit can be shared by many rings, limiting the total hardware overhead.

The rest of the paper is organized as follows: Section II describes the principles of on-chip optical communication. Section III discusses the variation in resonant wavelength induced by both process variation and dynamic thermal variation. Reliability challenges for optically connected systems are also discussed in this section. Our proposed channel remapping, sub-channel redundant rings, and self-calibration mechanism are illustrated in Section IV, followed by the simulation results in Section V. Section VI concludes the paper.

II. PRINCIPLES OF ON-CHIP OPTICAL COMMUNICATION

Ring resonators are widely used as electro-optical modulators, wavelength filters, and switches to redirect light from one waveguide to another. Ring resonators are constructed as circular waveguides placed adjacent to a signal waveguide, and each of these rings is sensitive to a certain set of resonant wavelengths. Multiple carrier signals are transmitted along the signal waveguide. If the resonant wavelength of a ring matches the wavelength of a carrier (shown as the solid curve in Figure 1), the resonator is considered to be on-resonance, and the carrier's light will be coupled into the ring, removing it from the waveguide.

As shown by the dashed curve in Figure 2, adding a bias voltage (V_B) to the ring will slightly shorten the resonant

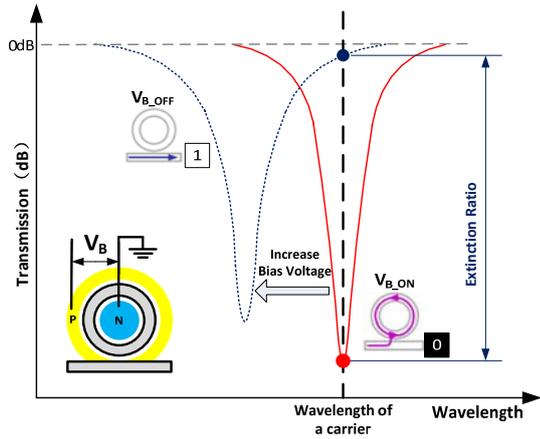


Figure 1. Transmission spectra of ring resonator w.r.t. bias voltage

wavelength, bringing the ring out of resonance with the carrier and allowing the light to pass.

Communication is performed by switching the bias voltage, bringing the resonator on-resonance (V_{B_ON}) and off-resonance (V_{B_OFF}) to directly modulate logic '0' and '1'.

Figure 2a shows the transmission (which is negatively proportional to resonance) spectrum of one ring resonator. Transmission minima appear repeatedly in the spectrum at each harmonic wavelength. The spacing between two such successive minima dictates the Free Spectral Range (FSR), which is the total available optical bandwidth.

Within the FSR, WDM (Wavelength Division Multiplexing) allows multiple light carriers to be transmitted along the same waveguide, significantly increasing its total data bandwidth. As shown in Figure 2b, λ_1 , λ_2 , and λ_3 are 3 different physical channels sharing one waveguide, whose wavelengths also fit within one FSR.

Figure 3 shows two example communication mechanisms using ring resonators. An off-chip laser light source generates multiple carriers, and injects them into a single signal waveguide. In Figure 3a, two carriers and corresponding modulator/detector pairs share one waveguide. R1 is biased with $V_B = V_{B_OFF}$, modulating '1' (letting light pass). The corresponding detector R3 couples this light in, and a Germanium doped photodiode embedded in the ring detects logic '1'. Similarly, R2 and R4 are using another channel to

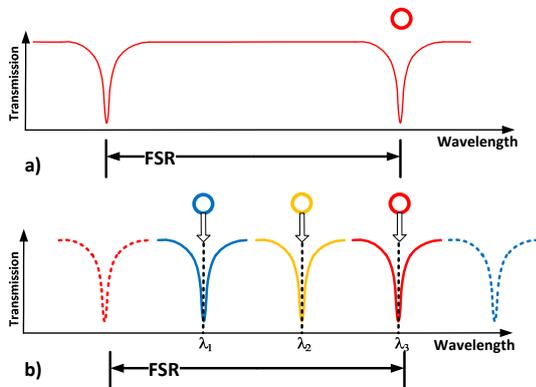


Figure 2. FSR and DWDM: a) FSR of one resonator; b) DWDM

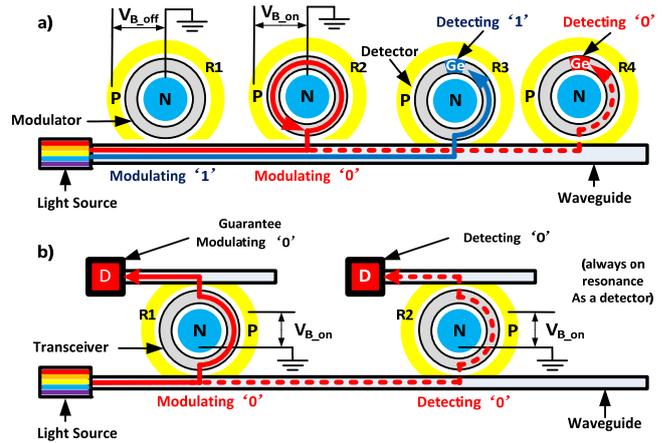


Figure 3. Mechanism of Electrical/Optical modulation and detection for one bit per channel a) using modulators and detectors; b) using transceivers

transmit logic '0': the carrier is absorbed by R2 and R4 senses the lack of light.

Figure 3b shows a different topology using transceivers rather than dedicated modulators and detectors. Although similar to a modulator, the transceiver has a drop waveguide with a detector at the end. When the ring is on resonance, light will be first coupled into the ring, then from there to the drop waveguide, and finally absorbed by the detector. Such a structure enables each ring to work as either a modulator or a detector. In the figure, transceiver R1 works as a modulator, modulating logic '0', and R2 detects '0' because no light is received by the detector at the end of its drop waveguide. Further, R1 can use its detector to monitor the quality of modulation. In this case, detector in R1 should detect '1' in order to guarantee R1's modulation of logic '0'.

III. RELIABILITY CHALLENGE

Ring resonators suffer from two kinds of variation: fabrication process variation and dynamic thermal variation [15], both of which cause the resonant wavelengths of the rings to shift from their design values. Combined, these variations present a reliability challenge for an optically connected system. It is possible to compensate for these variations by tuning the bias voltage and adding local heaters [16], but neither of these methods is sufficient to solve the problem within a reasonable power budget.

A. Sources of Variation

Process variation is deviation of the geometric parameters (ring radius, waveguide thickness, etc.) and will change the resonant wavelength, causing a mismatch between the resonant wavelength and the carrier wavelength. Previous studies have shown that a 1nm variation in ring radius will shift the resonant wavelength by approximately 0.58nm [15]. In [12], process variation caused as much as 1nm shift in resonant wavelength.

Thermal variations caused by power dissipation from the logic chip and surrounding environment will also cause shifts in resonant wavelength. Unlike process variations, which are fixed at manufacturing time, thermal variations change as the temperature of the environment and system workloads change. In [17], a 1°C change in temperature will lead to a shift of 0.11nm in resonant wavelength. According to [12], the

temperature could vary as much as 17°C between processor cores at any given time. For a system using resonators in [20] with FSR of 15nm and 64 channels on one waveguide, this amount of variation will lead to shifts as large as 8 channels. In the die-stacked structure required for an optically-connected system [11], the interposing layers between dies would work as a thermal low-pass filter. This will in turn make the spatial temperature distribution smoother in the optical die than the logic die and lower the speed of temperature change (1°C per second in [19]).

We used the parameters above to simulate the ring yield for optically connected systems like Corona. The simulation result shows that a mere 0.03 nm average shift in resonant wavelength will make 50% of the ring resonators inoperable.

B. Tuning methods

In order to build a functioning system, these variations require compensation in the form of tuning. There are two tuning methods commonly used to compensate for these variations: thermal tuning and bias voltage tuning.

Thermal Tuning (using local heaters): metal ohmic heaters are built on top of or beside the ring resonators, actively compensating for thermal and manufacturing variation. Such heaters are able to red-shift (i.e. lengthen) the resonant wavelength by 100 nm or more at a power cost of 2.4mW/nm [16] [14], 16fJ/bit [22], or lower by using athermal resonators [21]. Because of its large tuning range, local heating is necessary for a system with large variation, but local heating should be carefully applied because adding more heat to a microprocessor is not quite reasonable.

Bias voltage tuning: the bias voltage that shifts a ring modulator's resonant wavelength for modulation can also be adjusted to compensate for variation. Compared to thermal tuning, bias voltage tuning is relatively power efficient. In [25], only about 100uW is needed to blue-shift the resonant wavelength by 1nm. Unfortunately, the tuning range of bias voltage tuning is limited, and could be as little as 1nm [13]. As a direct tuning mechanism, we can use thermal tuning for large, coarse grained compensation, and then use bias voltage tuning for fine-grained control.

One supposed benefit of optical interconnects is lower power consumption. To realize this benefit, the total tuning power should be no more than 10W [23], which means on average 10uW per ring in a ~1 million ring system like Corona. With this tuning power budget, even the more power efficient bias voltage tuning can only compensate for an average shift of 0.1nm in resonant wavelength. Since this is far smaller than the expected variation, new approaches to tuning are needed.

IV. PROPOSED APPROACH

To overcome the obstacle discussed in the previous section, we propose a channel remapping with the use of sub-channel resonators to decrease the tuning power, and suggest a feedback mechanism to calibrate the bias voltages.

A. Channel Remapping

A 3-channel optical communication example is shown in Figure 4. Figure 4a shows 3 physical channels (i.e. carriers) λ_1, λ_2 and λ_3 evenly spaced within the free spectral range (FSR). The black dotted lines represent the center wavelengths of these carriers.

Ring resonators are fabricated for 3 logical channels LC1, LC2 and LC3, with nominal resonances at λ_1, λ_2 and λ_3 , respectively. The V-shaped curves represent their transmission spectra. Note the dotted curve in Figure 4a: this represents another resonance wavelength for the ring assigned to LC1. A similar dotted curve for LC3 is shown in Figure 4b and 4c. As we described in section II, a ring resonator will resonate at multiple harmonic wavelengths; the gap between them determines the FSR.

In Figure 4b, some combination of manufacturing and thermal variation has caused the resonance wavelengths of all three rings to red-shift by more than 1 channel, which requires significant tuning power to shift the resonant wavelengths of all rings back to their corresponding carriers. Due to the limited range of bias-voltage tuning and red-shift only nature of thermal tuning, the only option is to tune the next lower resonance peak for each ring (e.g. the dotted line for LC3) back to its respective carrier (the distance marked TRC).

On the other hand, if we remap the logical channels to different physical channels, we could tune the rings to the closest channel. This will limit the tuning distance (and in turn tuning power) to within one channel, even if the resonance deviation is far greater. An example of this is shown in Figure 4c, where we blue-shift a resonant wavelength of LC3 to the nearest carrier (λ_1) with bias voltage tuning. We can also tune the other two resonators to their closest carriers. Thus, we remap LC3 to λ_1 , LC1 to λ_2 and LC2 to λ_3 .

Consider this example with concrete numbers: For a ring resonator with an FSR of 62.5nm [25] attached to a waveguide with 64 carriers, the channel spacing is around 1nm. Without

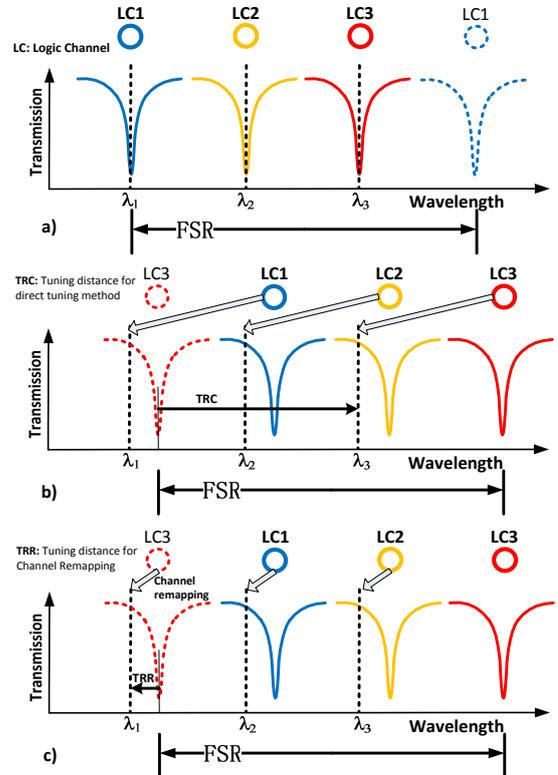


Figure 4. 3-channel optical communication example: a) Original design without variation; b) Compensating for large variation with direct tuning; c) Compensating with channel remapping

channel remapping, the maximum tuning range is an entire FSR (62.5nm), which will require significant thermal tuning. Channel remapping decreases this to about 1nm, allowing all the tuning to be done via bias voltage. A similar concept of using the FSR to remap the channels can be found in [26].

Recall that resonance deviation is caused by two forms of variation: process and thermal. Process variation is caused by factors such as etch depth and waveguide thickness. Most of these variations are spatially correlated: nearby rings tend to have approximately equal resonance shifts. The resonance variation within a cluster is within approximately 1nm [15].

Thermal variation is also spatially well-correlated. The 1.5 μ m-radius ring resonator in [25] has a footprint of about 25 μ m². For a 64-channel waveguide, 64 resonators clustered together will occupy less than 2000 μ m². Temperature across this area will be effectively equal, considering the 3-D chip thermal low-pass effect [19] and the thermal profile of a typical modern microchip [27].

We can therefore assume that the combined process and thermal variation will shift the resonant wavelengths of nearby rings the same amount (as implied in Fig. 4b and c). Under this assumption, channel remapping is simplified to cyclic shifting. Like in Figure 4b, after channel remapping, all logical channels are cyclically shifted one channel to the right.

While the amount of shift will be the same for adjacent rings (e.g. those assigned to one core), another group of rings elsewhere on the chip may shift a different amount. After tuning every ring to the nearest physical channel, we can send a pre-defined calibration packet. The cyclic shift required to realign this packet implies the shift needed for subsequent data packets received by that group of rings. Alternatively, the location of logical channel 1 could be encoded in the preamble of every packet.

There are two overhead costs associated with channel remapping: hardware and protocol. The hardware overhead should be minimal; a simple circular barrel shifter and some calibration sense circuitry are needed in the decode logic for each core. For an estimate of protocol overhead, consider a system with a communication speed of 10Gb/s [20], with calibration packets sent intermittently to determine the shift. The maximum rate of temperature change is 1°C/s [19], the resonant wavelength changes 0.11nm/°C [15], and channel spacing is 1nm. For a ring with -3dB bandwidth of 0.17nm [25], the worst case temperature change will take more than 773ms to shift the resonator out of resonance with the carrier. Therefore, the worst-case maximum spacing between calibration packets is 773ms. Compared to the 10Gb/s transmission speed, this is negligible.

B. Sub-Channel Redundant Rings

Channel remapping could reduce the average tuning range for each ring significantly, but still it might not satisfy the tuning power budget. In this subsection we propose a sub-channel redundant ring technique to further reduce the tuning power.

In Figure 5a, the rightmost ring must be tuned the distance indicated by TD. In Figure 5b, we add “half-channel” rings, one of which is shown as HCR, whose resonant wavelength is halfway between two adjacent original rings. For certain amounts of variation, these half-channel rings have a shorter tuning distance than the original rings (as shown in 5b). When

this occurs, we tune the half-channel rings to the carriers and use them instead of the original ones.

This technique can be generalized and extended beyond half-channel rings, as shown in Figure 5c, where 6 rings are available for each channel. As the number of redundant sub-channel rings increases, the tuning distance can be further reduced. This in turn reduces the tuning power.

One concern with additional redundant rings is area overhead. Based on [24], we anticipate the size of a processor chip to be approximately 260mm² in the near future. Recall that the optical devices are fabricated in their own die (in a 3D die-stacked system), so this entire area would be available. In a system with ~1 million rings operating on 64 channels like Corona, this leaves space for as many as 9 rings per channel.

Another concern with increasing numbers of redundant rings is decreased performance. Notice that in Figure 5c, after ring 2 is chosen for communication, adjacent rings 1 and 3 will still partly (and mistakenly) resonate with the carrier, adding undesired loss to the channel and potentially degrading communication. This problem will worsen as the number of sub-channel rings increases. To mitigate this, we need to tune the neighboring unselected rings away from resonance.

The process of using the redundant rings is as follows: 1) find the tuning distance of each redundant ring within one channel and select the one with shortest tuning distance; 2) tune the selected ring to the carrier and map its logical channel to that carrier; 3) tune away neighboring unselected rings if necessary. For spatially correlated variation, we only need to tune the rings of one channel and apply those results to all other rings in the same region.

Figure 6 is the simulation result of normalized tuning power with respect to the number of rings for one channel. We use the parameters in [25] for the simulation: FSR is 62.5nm shared by 64 carriers, channel gap is 1nm. Tune-away is performed when the center wavelengths of adjacent rings are close (within 2 times the -3dB bandwidth) to the operating ring.

For these parameters, we find that fabricating 5 rings per channel minimizes tuning power (reduced to only 22% of the original). With fewer rings, the increased tuning distance

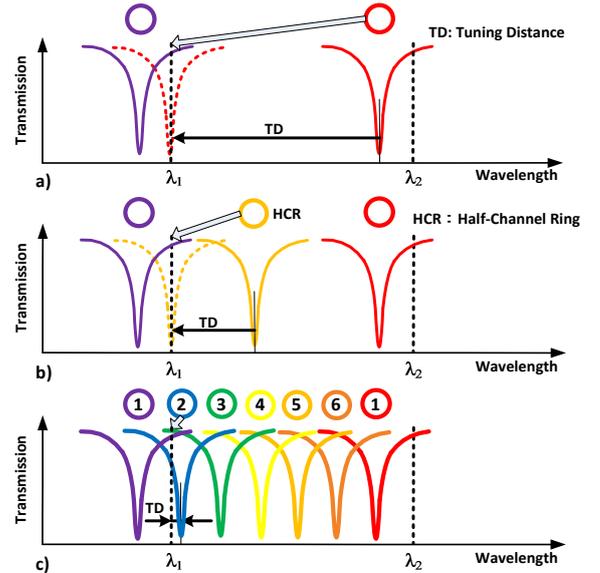


Figure 5. Examples of Sub-Channel Rings: a) original design; b) half-channel redundant; c) 1/6 channel redundant

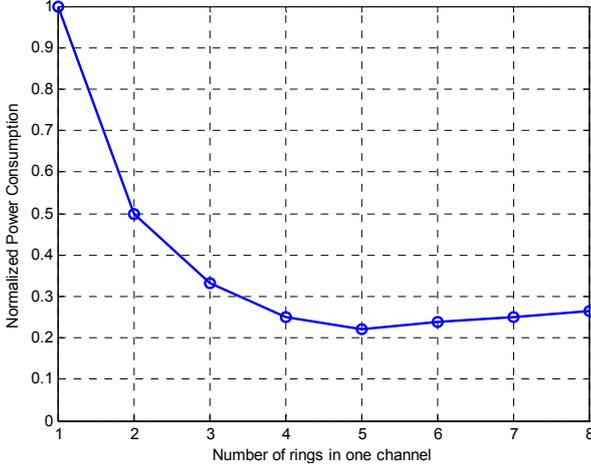


Figure6. Tuning power as a function of rings per channel

increases tuning power. With more rings, the tuning distance savings are outweighed by the additional power required to tune neighboring rings away.

C. Feedback Calibration Mechanism

A fundamental requirement for any calibration system, whether direct tuning or our proposed remapping scheme, is a dynamic feedback control circuit capable of constantly or periodically adjusting the tuning values to compensate for both process and temperature variation. In this subsection, we suggest a self-calibration circuit structure for transceiver-style ring modulator/detectors.

The basic calibration process is illustrated in Figure 7a. As the bias voltage V_B applied to the ring is swept through voltages V_{B1} , V_{B2} , and V_{B3} , it shifts the resonance spectra from C_1 to C_2 and C_3 . Figure 7b plots the modulation depth (ratio of light coupled in the ring resonator from the signal waveguide) during this sweep, with points P1, P2, and P3 corresponding with curves C_1 , C_2 , and C_3 in 7a. As shown, the modulation depth will first increase until maximal resonance is reached at P2, and then decrease.

An example tuning circuit structure is suggested in Figure 8. To begin calibration, the modulator is switched to “on-resonance”. The detector attached to the drop waveguide senses the coupled light intensity (and in turn the modulation depth), outputting a proportional voltage V_D . This voltage is fed into a maximum detector, which controls the bias sweep generator. When V_{D_MAX} is reached, the sweep stops. This corresponds to the maximal resonance point P2 in Figure 7.

The final output bias of the sweep is held as V_{B_ON} until the

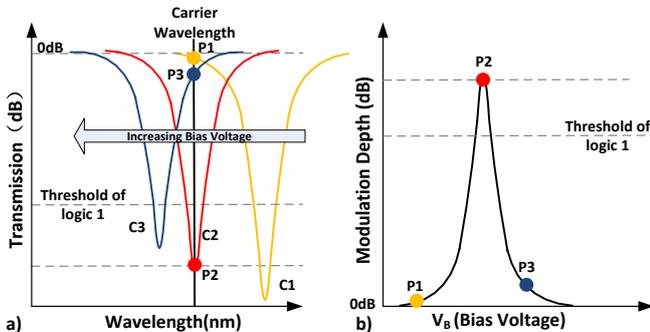


Figure7. Process of calibration

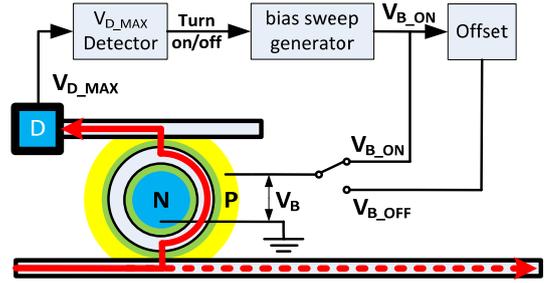


Figure8. Example of tuning circuit

next calibration cycle, and an offset is added to generate the off-resonance bias voltage V_{B_OFF} . An additional offset bias value (not shown) would also be needed to accommodate tune-away of neighboring sub-channel rings.

The area cost of this calibration circuit could be amortized across multiple rings through two sharing schemes. First, the spatial correlation of variation could allow one calibration circuit to determine the bias voltage for an entire temperature zone. Second, an entire set of sub-channel rings could share a single drop waveguide and detector without extra TSVs or mixed-signal circuitry, and be calibrated iteratively. Thresholds for minimum V_{D_MAX} and maximum V_B would ensure that only the on-resonance ring is selected.

V. SIMULATION RESULTS

To explore the potential power savings of our proposed technique, we simulate two values: 1) the tuning distance for one ring as a function of absolute temperature deviation (from nominal), and 2) the average tuning distance per ring in a system with different amounts of temperature variation across the chip. While total tuning power would be a preferred metric, it is difficult to model accurately without additional assumptions about heat dissipation and thermal conductivity.

Parameters used in the simulations are as follows: channel gap is 1nm [11][25], average variation in resonant wavelength by process variation is 1nm [15], and temperature deviation will cause 0.11nm/°C [12] shift in resonant wavelength. The first simulation is run both with and without process variation to highlight the effects of the proposed technique.

In Figure 9, using direct tuning with no remapping, the tuning distance increases with temperature deviation. As the resonance peak is shifted farther and farther away from the carrier wavelength, more power is needed to tune it back.

With channel remapping, the tuning distance will periodically return to 0 every ~ 9 degrees (ignoring process variation): at this point, the channel remapping has simply shifted the logical channel to the next physical carrier, with which it is precisely aligned. The use of 5 (sub-) channel rings exhibits the same behavior, but with a reduced period corresponding to 1/5 of the channel spacing.

Process variation is assumed to be a random variable with mean value of 0, and in some cases will partially cancel out the effect of thermal variation. As shown in Figure 9, this has the effect of smoothing out the tuning distance to approach its average. Additional process variation will merely cause this smoothing effect to be more pronounced.

In Figure 10, we simulate the average per-ring tuning distance across an entire optically connected chip as a function of the temperature difference between two temperature zones. The temperature of each zone is a random variable with mean 0

(i.e. nominal), and the maximum temperature variation is expressed in terms of 6 times the standard deviation.

As shown in the figure, when max temperature variation reaches 17°C [15], the tuning distance for direct tuning could be as much as 4.60nm. Note also that for temperature variation exceeding about 1°C, the average tuning distance for channel remapping and sub-channel rings levels off at only 0.44 nm and 0.097 nm (2.1%), respectively. This means that with a combination of channel remapping and sub-channel rings, the total tuning power with bias voltage tuning for a system with 1 million rings [10] could feasibly be brought under 10W [23].

V. CONCLUSION

In this paper, we propose a dynamic reconfiguration and calibration technique that helps to enable on-chip optical communications by significantly reducing the required tuning power. We also propose a self-calibration mechanism that would work in tandem with these techniques. Simulation results show that using the proposed approach, the average tuning distance could be decreased to about 2.1% of the original, bringing the total tuning power under 10W.

ACKNOWLEDGEMENT

This research is partially supported by the Gigascale Systems Research Center (GSRC), one of six research centers

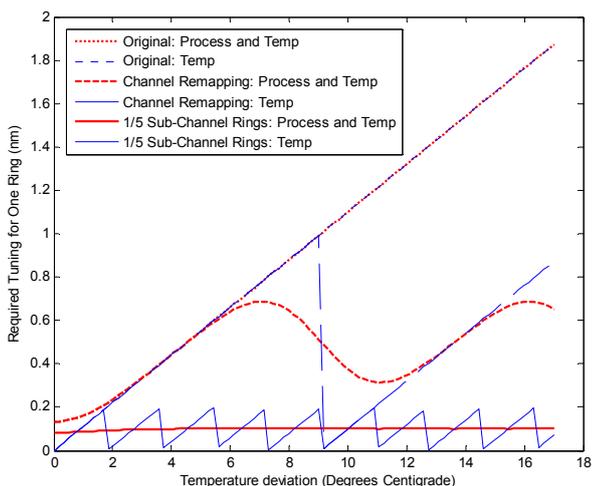


Figure9. Required Tuning for one ring WRT Temp deviation

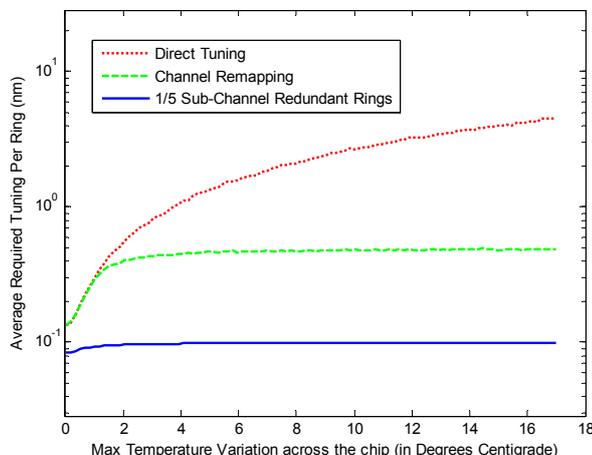


Figure10. Average required tuning of one ring WRT max temp variation

funded under the Focus Center Research Program and supported by the Semiconductor Research Corporation.

REFERENCE

- [1] Ian A. Young, et al., "Optical I/O Technology for Tera-Scale Computing" IEEE J. Solid-State Circuits, vol 45, no. 1, pp 345-248, 2010
- [2] G. T. Reed, et al., "Silicon optical modulators," Nature Photonics, July 2010.
- [3] Y. Shen et al., "Nanoscale oxidation of silicon microring resonators," in CLEO, 2011
- [4] Q. Xu et al., "12.5 Gbit/s carrier-injection-based silicon micro-ring silicon modulators," Optics Express, vol. 15, no. 2, pp. 430-436, 2007.
- [5] J. K. S. Poon et al., "Wide-range tuning of polymer microring resonators by the photobleaching of the CLD-1 chromophores," Optical Letters, vol. 29, no. 22, pp. 2584-2586, 2004.
- [6] M. R. Watts et al., "Ultralow power silicon microdisk modulators and switches," in IEEE Int. Conf. on Group IV Photonics, pp. 4-6, Sep. 2008.
- [7] D. Liang, et al., "Native-oxide-confined high-index-contrast lambda=1.15 um strain-compensated InGaAs single quantum well ridge waveguide lasers," Applied Physics Letters, vol, 93, no. 16, pp. 161-108, 2008.
- [8] A. Shacham, et al., "Photonic Networks-on-Chip for Future Generations of Chip Multiprocessors", Trans. Computer, Vol. 57, No. 9, pp:1246-1260
- [9] N. Sherwood-Droz et al., "Scalable 3D dense integration of photonics on bulk silicon," Opt. Express, vol. 19, pp. 17758-17765, 2011
- [10] J. Ahn, et al., "Devices and architectures for photonic chip-scale integration" Appl Phys A, pp: 989-997, 2009
- [11] D. Vantrease et al., "Corona: system implications of emerging nanophotonic technology," in IEEE International Symposium on Computer Architecture, pp. 153-164, 2008.
- [12] Li Z., et al. "Reliability modeling and management of nanophotonic on-chip networks," IEEE Trans. VLSI Systems, no. 99, pp. 1-14, 2011.
- [13] Q. Xu et al., "12.5 Gbit/s carrier-injection-based silicon micro-ring silicon modulators," Optics Express, vol. 15, no. 2, pp. 430-436, 2007.
- [14] M. A. Popovic, "Theory and design of high-index-contrast microphotonic circuits," Ph.D. dissertation, MIT, Cambridge, 2008.
- [15] L. Zheng, et al., "Device Modeling and System Simulation of Nanophotonic on-Chip Networks for Reliability, Power and Performance", in DAC, Jun. 2011
- [16] Po Dong, et al. "Broadly tunable high speed silicon micro-ring modulator," in IEEE Photonics Society Summer Topical Meeting Series, pp. 197-198, July 2010.
- [17] S. Manipatruni et al., "Wide temperature range operation of micrometerscale silicon electro-optic modulators," Optical Letters, vol. 33, no. 19, pp. 2185-2187, 2008.
- [18] A. Krishnamoorthy, "Low-power, high-density optical interconnects to the processor", in OFC, March, 2011
- [19] Li Z. "On-chip optical and electrical communication: analysis and design for system level performance and reliability," Ph.D. dissertation, Tsinghua University, Beijing, 2010.
- [20] Q. Xu et al., "Micrometerscale silicon electro-optic modulator" Nature, vol. 435, pp. 325-327, May 2005.
- [21] Linjie Zhou, et al. "Towards athermal optically-interconnected computing system using slotted silicon microring resonators and RF-photon comb generation", Appl Phys A, vol 95, pp:1101-1109, 2009
- [22] Ajay Joshi, et al. "Silicon-Photonic Clos Networks for Global On-Chip Communication", in NOCS-3, May 2009
- [23] Z. Li et al., "Global on-chip coordination at light speed," IEEE Design & Test of Computers, vol. 27, no. 4, pp. 54-67, July 2010.
- [24] International Technology Roadmap for Semiconductors; <http://www.itrs.net/>, 2010.
- [25] Q. Xu, D. Fattal, and R. G. Beausoleil, "Silicon microring resonators with 1.5-um radius," Optical Express vol. 16, pp. 4309-4315, 2008.
- [26] Nathan Binkert, et al., "The Role of Optics in Future High Radix Switch Design" in ISCA'11, pp:437-447, 2011.
- [27] Marty M R, et al., "Virtual hierarchies to support server consolidation," in ISCA07, pp: 46-56, 2007.