System-Level Modeling and Analysis of Thermal Effects in WDM-Based Optical Networks-on-Chip

Yaoyao Ye, *Member, IEEE,* Zhehui Wang, *Student Member, IEEE,* Peng Yang, *Student Member, IEEE,* Jiang Xu, *Member, IEEE,* Xiaowen Wu, *Student Member, IEEE,* Xuan Wang, *Student Member, IEEE,* Mahdi Nikdast, *Member, IEEE,* Zhe Wang, *Student Member, IEEE,* and Luan H. K. Duong, *Student Member, IEEE*

***Abstract*—Multiprocessor systems-on-chip show a trend toward integration of tens and hundreds of processor cores on a single chip. With the development of silicon photonics for short-haul optical communication, wavelength division multiplexing (WDM)-based optical networks-on-chip (ONoCs) are emerging on-chip communication architectures that can potentially offer high bandwidth and power efficiency. Thermal sensitivity of photonic devices is one of the main concerns about the on-chip optical interconnects. We systematically modeled thermal effects in optical links in WDM-based ONoCs. Based on the proposed thermal models, we developed OTemp, an optical thermal effect modeling platform for optical links in both WDM-based ONoCs and single-wavelength ONoCs. OTemp can be used to simulate the power consumption as well as optical power loss for optical links under temperature variations. We use case studies to quantitatively analyze the worst-case power consumption for one wavelength in an eight-wavelength WDM-based optical link under different configurations of low-temperature-dependence techniques. Results show that the worst-case power consumption increases dramatically with on-chip temperature variations. Thermal-based adjustment and optimal device settings can help reduce power consumption under temperature variations. Assume that off-chip vertical-cavity surface-emitting lasers are used as the laser source with WDM channel spacing of 1 nm, if we use thermal-based adjustment with guard rings for channel remapping, the worst-case total power consumption is 6.7 pJ/bit under the maximum temperature variation of 60 °C; larger channel spacing would result in a larger worst-case power consumption in this case. If we use thermal-based adjustment without channel remapping, the worst-case total power consumption is around 9.8 pJ/bit under the maximum temperature variation of 60 °C; in this case, the worst-case power consumption would benefit from a larger channel spacing.**

***Index Terms*—Multiprocessor, optical interconnect, optical network-on-chip, temperature sensitivity, thermal effect, WDM.**

I. Introduction

MULTIPROCESSOR systems-on-chip (MPSoCs) show a trend toward integration of tens and even hundreds of processor cores on a single chip. The growing on-chip communication demands put tremendous pressure on the communication architecture, and it has become a limitation to the continuous performance improvement of multiprocessor systems. Networks-on-chip (NoCs) are efficient and scalable on-chip communication architectures for MPSoCs [1], [2]. Instead of routing design-specific global interconnects, information is exchanged by routing packets in the network based on modern networking theories. With technology scaling, the on-chip communication demands an increasing proportion of the system power budget. In current prototypes with tens of cores, the power consumed by the electronic NoC accounts for over 25% of the overall power and this is too high to meet the expected requirements of future multiprocessor systems [3]. The use of optical interconnects can help manage the power budget in multiprocessor architectures [4]. Silicon photonics-based optical on-chip interconnects offer fundamental physical advantages to overcome the limitations faced by metallic interconnects.

Optical networks-on-chip (ONoCs) are proposed as an emerging communication architecture for new-generation MPSoCs. ONoCs can potentially offer ultrahigh communication bandwidth, low latency, and high power efficiency to on-chip communications. Most ONoC architectures employ silicon photonic devices which can be integrated with existing CMOS-based processor cores either through CMOS-compatible fabrication processes or bonding technologies. Recent developments in nanoscale silicon photonic devices substantially improve the feasibility of ONoCs [5]. Microresonator (MR) based add-drop filters have been widely used as the basic optical switching element (BOSE) to perform the switching functions in ONoCs [6]–[8]. Optical transmitter converts electrical signals into optical signals by using an optical modulator or directly modulating the driving current of a vertical-cavity surface-emitting laser (VCSEL) [9]. VCSEL is an attractive candidate for on-chip laser source because of its low power consumption, manufacturing advantages and high modulation bandwidth at 10–40 Gb/s. 3-D integration technology can be used to connect VCSELs with the underlying CMOS driver circuits through silicon via (TSVs). Optical receiver uses photodetectors to convert optical signals into electrical signals used by processors or memories in electronic domains. A typical receiver also includes transimpedance amplifier (TIA) and limiting amplifier (LA) circuits for current-to-voltage conversion and voltage amplification.

Manuscript received September 7, 2013; revised January 30, 2014 and April 22, 2014; accepted July 15, 2014. Date of current version October 16, 2014. This work was supported under Grant GRF620911, Grant GRF620512, and Grant DAG11EG05S. This paper was recommended by Associate Editor R. O. Topaloglu.

The authors are with the Department of Electronic and Computer Engineering, Hong Kong University of Science and Technology, Hong Kong (e-mail: yeyaoyao@connect.ust.hk; jiang.xu@ust.hk).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TCAD.2014.2351584

ONoCs take advantage of silicon photonics and NoC architectures to enhance the performance and power efficiency for on-chip communication. However, as an intrinsic characteristic of photonic devices, thermal sensitivity is one of the major concerns in ONoC designs. Due to the nonuniform power densities across the chip as well as the limited thermal conductivity of the die and packaging materials, steady-state chip temperature varies spatially. In general, chip temperatures can vary by more than 30 °C across the chip under typical operating conditions [10]. As the number of cores on the chip is increasing to tens of cores and even hundreds of cores, the on-chip temperature gradients could be even higher [11]. As a result of the thermo-optic effect, the characteristics of silicon photonic devices change with temperature variations. Optical link components such as laser source and BOSE are sensitive to temperature variations. Such thermal effects can potentially cause additional optical power loss in ONoCs. Without thermal variations, the projected power efficiency of ONoCs is on the order of 1 pJ/bit. However, if with thermal variations, the extra thermal-induced power loss would degrade the power efficiency. The temperature-dependent wavelength shift in a MR-based BOSE would result in more insertion loss at the drop port, especially for a high-Q MR. For example, for a MR in 1550 nm with quality factor Q of 10^4, the 3 dB bandwidth is 0.155 nm, a temperature variation of 3 °C can make the spectrum shift by one 3 dB bandwidth. To compensate for a high optical power loss caused by chip temperature variations, more input power would be needed by the transmitter to guarantee enough optical power reaching the receiver.

In this paper, we systematically model the thermal effects in WDM-based optical links, and develop a system-level analytical thermal model for WDM-based ONoCs. The work involves the modeling of optical power loss of WDM-based optical link components under temperature variations, including the BOSE, the basic optical modulation element (BOME), and the basic optical filter element (BOFE). Based on the thermal models, we developed OTemp, an optical thermal effect modeling platform for both WDM-based optical links and single-wavelength optical links. OTemp can be used to simulate the power consumption as well as optical power loss of optical links under temperature variations. We use case studies to show the worst-case power consumption of WDM-based ONoCs under temperature variations. In addition, we analyze the trade-offs in choosing device settings to reduce the thermal-induced power consumption, involving the setting of WDM channel spacing, the setting of the basic optical components, the use of thermal-based adjustment to compensate for thermal variations, etc. In addition to thermal variations, process variation is another source of the wavelength variations of microresonators [12]. In this paper, we focus on the thermal effects caused by temperature variations. We will include more analysis for process variations in future work.

The rest of paper is organized as follows. Section II gives a survey of related works on ONoC architecture designs and the investigations of thermal sensitivity. Section III presents the systematic thermal modeling for a WDM-based optical link. Section IV introduces OTemp, an optical thermal effect modeling platform for both WDM-based optical links and single-wavelength optical links. In Section V, we use case studies to quantitatively study the worst-case power consumption of WDM-based ONoCs under different configurations of low-temperature-dependence techniques. Section VI draws the conclusions of this paper.

II. Related Work

With the booming developments in nanoscale silicon photonic technologies, ONoCs with various architectures have been proposed in literature. Shacham *et al.* [6] proposed an augmented-torus ONoC based on 4 × 4 optical switches. Mo *et al.* [13] proposed a hierarchical mesh-based ONoC. Gu *et al.* [7] proposed a fat-tree based ONoC. Ding *et al.* [14] presented an optical routing framework to reduce the power consumption of ONoCs with the technique of integer linear programming. WDM technology can enable much higher optical bandwidth of optical interconnects. Lots of related works explore the ultrahigh performance of ONoCs with the support of WDM technology. Vantrease *et al.* [15] proposed a clustered all-optical interconnection network, Corona, with dense WDM support. Each channel consists of four waveguides, and each waveguide is with 64 wavelengths. Pan *et al.* [8] proposed the dense-WDM based Firefly architecture, which is an optical crossbar-based architecture with localized arbitration. This paper assumed to use an off-chip laser source which provides 64 wavelengths for the WDM-based optical channel [8]. Kirman *et al.* [16] proposed a WDM-based hierarchical optical bus for multiprocessor systems. A range of 4 to 12 wavelengths per waveguide was explored for the proposed on-chip optical bus. Briere *et al.* [17] presented a multistage ONoC with a passive-switching router. The work couples the wavelength routing method to the WDM technique. Beausoleil *et al.* [18] proposed a crossbar-based all-optical arbitration network with dense WDM. Pasricha and Dutt [19] presented an ONoC based on an optical ring with bus protocol standards. WDM is used to improve the bandwidth density of the proposed optical interconnect. Cianchetti *et al.* [20] proposed a WDM-based optical network with low-latency predecoded source routing. Kodi *et al.* [21] proposed an ONoC architecture with static routing and wavelength allocation, which involves selective merging of multiple wavelengths to maximize the bandwidth. Batten *et al.* [22] proposed an optical mesh based on a hybrid optical-electrical global crossbar, where processing cores and DRAM are divided into sub-mesh and connected with the optical crossbar. The proposed architecture supports dense WDM for advantages of performance and energy efficiency.

Though ONoCs offer a new approach to empowering ultrahigh bandwidth with low power consumption, there are concerns about the reliability of optical interconnects for on-chip applications. An investigation of thermal issues of on-chip optical interconnects shows that, with the consideration of thermal regulation power, optical interconnects may not have advantages in power efficiency as compared with their electrical counterpart [23]. Previous works have investigated the thermal sensitivity of silicon photonic devices. The temperature-dependent wavelength shift in silicon-based

microresonators is found to be about 50–100 pm/K, which is nonnegligible in practical use [24]–[26]. Thermal-based adjustment by local microheaters is an alternative solution to compensating for the temperature-dependent wavelength shift of microresonators and the tuning efficiency in current technology is on the order of several mW/nm [27], [28]. Another possible solution is to fabricate microresonators with low temperature dependence [29]. Besides microresonators, lasers are also temperature sensitive. Investigations of the temperature sensitivity of VCSELs show that the temperature-dependent wavelength shift is comparable to or even larger than that of microresonators [30], [31]. Furthermore, because of the mutual shift between lasing wavelength and peak material gain wavelength under temperature variations, VCSEL power efficiency degrades seriously at high temperatures [9]. There are also investigations of the temperature-dependent behaviors of Ge-based photodetectors [32], [33].

As a result of thermo-optic effect, on-chip temperature variations can potentially cause ONoC performance degradation and much more power consumption. An ONoC thermal model is required to fully understand these challenges and help further develop this emerging technology. In [34], we proposed the thermal modeling for single-wavelength based optical links. The thermal analysis is based on an abstracted single-wavelength optical link model which is composed of a directly-modulated laser, single-ring optical switching elements, and a photodetector. Based on the single-wavelength ONoC thermal model, we revealed three important factors regarding ONoC power efficiency under temperature variations, including the initial setting of photonic devices, the number of switching stages in the ONoC architecture, and the bandwidth of optical switching elements. ONoCs with WDM support can achieve much higher performance by combining multiple signals at different wavelengths for transmission. As compared to the single-wavelength ONoCs, WDM-based ONoCs are based on different optical components which supports multiple wavelengths, such as multiple-ring optical switching element, multiple-ring optical modulation element and multiring optical filter element at the receiver. As a result, the thermal models for single-wavelength ONoCs [34] cannot be applied to WDM-based optical links. A new system-level thermal model is required to be developed for WDM-based ONoCs consisting of multiwavelength optical components.

This paper presents new thermal models and analysis respectively for the multiple-ring optical switching element, the multiple-ring optical modulation element and the multiple-ring optical filter element. Based on the new thermal models for each component, this paper presents the link-level thermal model which reveals the relationship among all the components. The major contribution of this paper is that we provide a general model to study thermal effects in WDM-based ONoCs with different configurations (e.g., the number of wavelengths in WDM, channel spacing, off-chip or on-chip laser, thermal based adjustment, direct modulation or using separate modulators, the initial setting of microresonators, electronic-based or thermal based switching, the number of active switching stages, the number of active switching stages, etc.). This paper is not about how to obtain temperature map (or how the heat

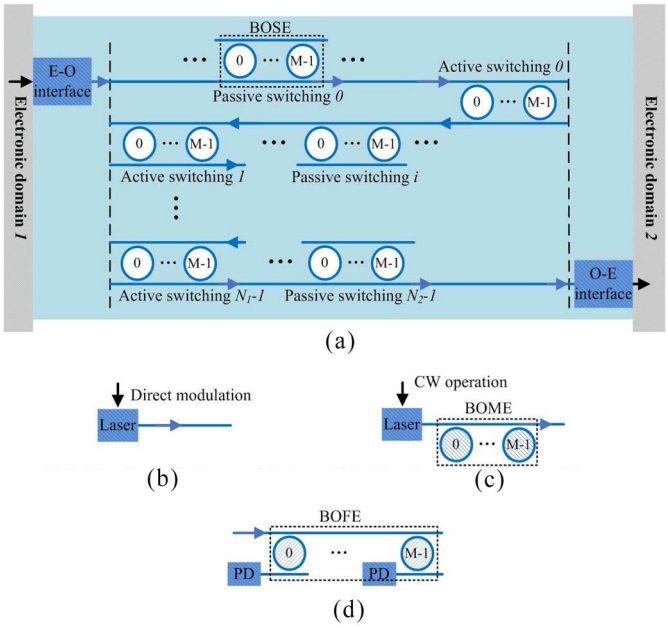

Fig. 1. (a) Overview of an M-wavelength WDM-based optical link in ONoCs connecting two electronic domains. (b) and (c) E-O interface with direct modulation and separate modulators. (d) O-E interface.

dissipated by circuitry affects chip temperature), but is focused on how temperature variations affect the characteristics of WDM based ONoCs. The proposed thermal model is general enough that can be applied to any thermal maps. We not only model the thermal effects in all the individual components in a WDM based optical link, but also find out a mathematical formula which reveals the mathematical relationship among the components. To our knowledge, no one has proposed such formula in published literature. The silicon photonics based optical on-chip interconnect is a new research area, and is still in the initial stage. It costs very much for a hardware prototype implementation and measurement. Due to the high cost and the limited testing techniques, most works and simulators in this research area rely on theoretical reasoning and analysis. We also choose the theoretical reasoning and analysis as the approach. In our modeling, well-known fundamental device-level models which have been experimentally verified serve as the solid foundation. Our work is a system-level model which puts together the experimentally-demonstrated device-level basic models. The proposed model is reliable as long as the theoretical framework is unflawed.

III. System-Level Thermal Modeling of WDM-Based Optical NoCs

Fig. 1 shows the general overview of an M-wavelength WDM-based optical link in ONoCs. Despite the architecture diversity, an optical link in WDM-based ONoCs is generally composed of an E-O interface (optical transmitter), an optical path, and an O-E interface (optical receiver). The E-O interface converts electrical signals into M-wavelength WDM optical signals by directly modulating the driving current of the WDM laser source [Fig. 1(b)], or using the BOME [Fig. 1(c)]. On the optical path between the optical transmitter and optical receiver, the BOSEs are used to switch optical signals until

they reach the destination. If the input optical signals are on-resonance with the BOSE, they would be coupled into the BOSE and be directed to the drop port (active switching); otherwise the optical signals would pass through the BOSE without changing the propagation direction (passive switching). In the optical link model, we assume the number of active switching and passive switching is N_1 and N_2 respectively. At the receiver side, the O-E interface includes photodetectors (PD) which convert optical signals into electrical signals used by processors or memories in electronic domains. MR-based BOFE [Fig. 1(d)] is used at the receiver to demultiplex the M wavelengths before the optical signals reaching the photodetectors.

A. Thermo-Optic Effect

As a result of the thermo-optic effect, material refractive index is temperature dependent and follows (1), where n_0 is the refractive index at room temperature, dn/dT is the thermo-optic coefficient of the material, and ΔT is the temperature variation. Physical measurements show that the thermo-optic coefficient of silicon is on the order of $10^{-4}/K$ and is nonlinear over a large temperature range at 1550 nm wavelength [35]

$$n = n_0 + \frac{dn}{dT}\Delta T. \tag{1}$$

Since the refractive index is an important device parameter, the thermo-optic effect will cause changes in the device characteristics. In relation to the MR, the resonance is directly governed by the effective index. The resonant wavelength of microresonators red-shift approximately linearly with increasing temperatures. For a MR working at temperature T_{MR}, the resonant wavelength is as in (2), where $\lambda_{\mathrm{MR\_0}}$ is the resonant wavelength at room temperature T_0, and ρ_{MR} is the temperature-dependent wavelength shift coefficient of the MR

$$\lambda_{\mathrm{MR}} = \lambda_{\mathrm{MR\_0}} + \rho_{\mathrm{MR}}\,(T_{\mathrm{MR}} - T_0). \tag{2}$$

The lasing wavelength of VCSEL also red-shifts approximately linearly with temperature. If using off-chip VCSEL as the laser source, the lasing wavelength can be fixed by equipping with a temperature control unit. If using on-chip VCSEL as the laser source, the temperature-dependent wavelength shift and power efficiency degradation should be taken into account in the thermal model. The output power of VCSEL degrades at higher operating temperatures. Assuming that the VCSEL is driven by current I which is above the threshold but before the point where the output power starts to decrease with the current, we can express the output optical power P_{out} by (3), where T_{VCSEL} is the VCSEL operating temperature, I is the driving current, α is the minimum threshold current, T_{th} is the temperature at which the threshold current is the minimum, β is a coefficient related to the temperature dependance of the threshold current, ε is the slope efficiency at 0 °C, and γ is a positive coefficient. For the VCSEL demonstrated in [9], when the temperature changes from room temperature to 80 °C, the slope efficiency decreases from 0.36 to 0.23 mW/mA, and the maximum emission power decreases from 4 to 1.5 mW correspondingly

$$P_{\mathrm{out}} = \left(I - \alpha - \beta\,(T_{\mathrm{VCSEL}} - T_{\mathrm{th}})^2\right)(\varepsilon - \gamma \cdot T_{\mathrm{VCSEL}}). \tag{3}$$

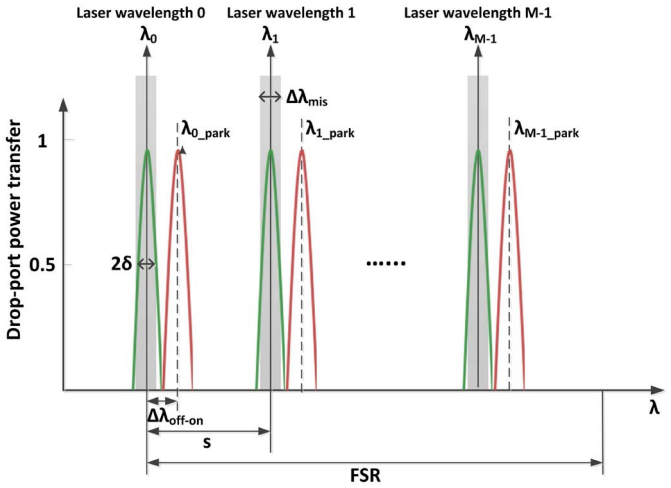

Fig. 2. Electronic-controlled BOSE baseline setting.

B. Thermal Model of the BOSE

Fig. 2 shows the baseline setting of an electronic-controlled M-wavelength BOSE. We assume the WDM laser sources provide M wavelengths at $\lambda_0, \lambda_1, \ldots, \lambda_{M-1}$ (at room temperature), with a channel spacing of s. When the BOSE is in the off state, the microresonators park away from the corresponding laser wavelengths by $\Delta\lambda_{\mathrm{off-on}}$. When the BOSE is turned on by applying forward-bias voltage, the microresonators blue-shift to the corresponding laser wavelengths $\lambda_0, \lambda_1, \ldots, \lambda_{M-1}$, respectively. On-resonance WDM optical signals would be actively switched to the drop port of the BOSE with a small insertion loss.

The overall amplitude transfer at the drop port of an M-wavelength BOSE can be found by application of the recursive formula in (4), where the recursion starts with $f_0 = r_0$ [36]. The insertion loss of an active BOSE is as (8). If the MR is turned on (blue-shift) by carrier injection, there would be an extra absorption loss. ω_n is the on-state resonance frequency of the ring n. For BOSEs with temperature variation ΔT, the on-state resonance frequency of ring n would shift away from the corresponding laser wavelength with the temperature-dependent wavelength shift of ρ_{MR} (5). More insertion loss in the active BOSE would be caused by the temperature-dependent wavelength shift. The wavelength shift caused by process variations can be directly added to the wavelength shift modeled in (5). r_n (6) and t_n^i (7) are the amplitude transmission from the input port to the drop port and through port respectively for ring n in isolation ($n = 0, 1, \ldots, M-1$) [37]. t_n^o is the amplitude transmission from the add port to the drop port. $\theta_n = 2\pi L_n/\lambda$ is the phase delay along the waveguide between the reference planes of ring n and ring $n+1$, where L_n is the spacing between the two rings and λ is the wavelength in free space. $\tau = 2\pi R n_g/c$ is the round-trip propagation time in MR. We assume that microresonators in the array are coupled to waveguides with coupling coefficient of κ^2

$$f_n = r_n - \frac{t_n^i t_n^o}{r_n - f_{n-1}^{-1} exp\,(j2\theta_{n-1})} \tag{4}$$

$$\omega_n = \frac{2\pi c}{\lambda_n + \rho_{\mathrm{MR}} \cdot \Delta T} \tag{5}$$

$$r_n = \frac{2\kappa^2}{j2\tau\,(\omega-\omega_n)+\left(2\kappa^2+\kappa_p^2\right)} \tag{6}$$

$$t_n^i = t_n^o = \frac{j2\tau\,(\omega-\omega_n)+\kappa_p^2}{j2\tau\,(\omega-\omega_n)+\left(2\kappa^2+\kappa_p^2\right)} \tag{7}$$

$$L_{\text{BOSE\_active}} = -10\log|f_{M-1}|^2 \tag{8}$$

$$L_{\text{BOSE\_parking}} = -10\log\left(1-|f_{M-1}|^2\right). \tag{9}$$

As shown in Fig. 2, if without temperature variation, each ring in a passive BOSE (the red curves) is parked far away from the laser wavelengths. Since the laser wavelengths are out of the resonance with the rings in the passive BOSE, it will only insert a small loss to passing optical signal. However, parking BOSEs under temperature variations may wrongly shift into the misplace region (defined as the gray bars in Fig. 2) of the neighboring laser wavelengths. $\Delta\lambda_{\text{mis}}$ is defined as the width of the misplace region of each laser wavelength. For a parking BOSE red-shifting by $\rho_{\text{MR}}\cdot\Delta T$, if the condition in (10) is satisfied, it is considered to be misplaced at the laser wavelength λ_x ($x\in[0,M-1]$)

$$\Delta\lambda_{\text{off-on}}+\rho_{\text{MR}}\cdot\Delta T-s\cdot x \in[-0.5\cdot\Delta\lambda_{\text{mis}},0.5\cdot\Delta\lambda_{\text{mis}}]. \tag{10}$$

When the rings resonance of the parking BOSE shifts into the misplace region of each laser wavelength, since the laser wavelengths are close to rings resonance, part of the power will be coupled into the rings and this would introduce significant insertion loss to passing laser wavelengths. According to (9), if $\Delta\lambda_{\text{mis}}$ is set as three times of the MR 3 dB bandwidth 2δ, the insertion loss of the parking BOSE is 0.46 dB when the rings resonance in the parking BOSE are at the edge of the misplace region of each laser wavelength; if $\Delta\lambda_{\text{mis}}$ is set as 2δ, the corresponding insertion loss is 3 dB. With a predefined $\Delta\lambda_{\text{mis}}$, thermal-based adjustment can be used to move the parking BOSE out of the misplace region. In this mechanism, the thermal-based adjustment will not be activated until when the parking BOSE shifts into the misplace region. By setting a small $\Delta\lambda_{\text{mis}}$, the insertion loss of the parking BOSE may be quite significant before the thermal-based adjustment is activated. By setting a large $\Delta\lambda_{\text{mis}}$, the thermal-based adjustment will be timely applied to the parking BOSE to avoid significant insertion loss. For example, if we conduct the thermal-based adjustment for parking BOSEs with $\Delta\lambda_{\text{mis}}$ of three time of 2δ, the worst-case insertion loss of the parking BOSE is 0.46 dB before the thermal-based adjustment is activated.

Fig. 3 shows the insertion loss of an eight-wavelength parking BOSE with channel spacing of 1, 2.355, and 4.155 nm, respectively. We assume the switching of BOSE is electronic-controlled with $\Delta\lambda_{\text{off-on}}$ of 0.4 nm, and the qualify factor Q of microresonators is 5000. With channel spacing of 1 nm (the baseline setting), the insertion loss of a parking eight-wavelength BOSE could be as high as 10 dB when the parking BOSE is inside the misplace region (e.g., with temperature variation of 10 °C). In order to prevent the thermal-induced misplacement of parking BOSE without online thermal-based adjustment, the WDM channel spacing s should be set as large

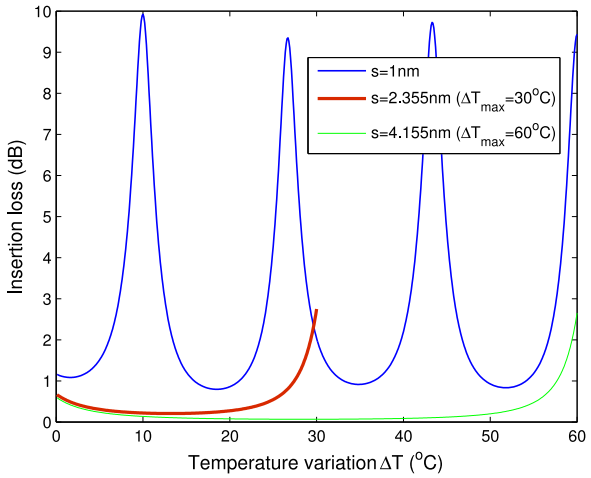

Fig. 3. Insertion loss of a parking eight-wavelength BOSE with temperature variations (electronic-controlled switching).

as (11), where ΔT_{max} is the maximum temperature variation of the parking BOSE

$$s\geq\Delta\lambda_{\text{off-on}}+\rho_{\text{MR}}\cdot\Delta T_{\text{max}}+0.5\cdot\Delta\lambda_{\text{mis}}. \tag{11}$$

According to (11), the channel spacing of 2.355 nm is chosen to guarantee that the worst-case insertion loss of the parking BOSE before activating the thermal-based adjustment is under 3 dB for a maximum temperature variation of 30 °C. Thus the red curve in Fig. 3 is plotted only until to a temperature variation of 30 °C. The channel spacing should be no less than 2.665 nm in order to guarantee that the worst-case insertion loss of the parking BOSE before activating the thermal-based adjustment is no larger than 0.46 dB. For ΔT_{max} of 60 °C, if the channel spacing is as large as 4.155 nm, the insertion loss of the parking BOSE before activating the thermal-based adjustment would be no larger than 3 dB at the worst case. The channel spacing should be no less than 4.465 nm to guarantee that the insertion loss of the parking BOSE before activating the thermal-based adjustment would be no larger than 0.46 dB at the worst case.

By using a large channel spacing, although it prevents the misplacement of parking BOSEs, it may result in more insertion loss in active BOSEs. The power transfer spectral of an active BOSE has multiple peaks, where each peak corresponds to the resonance of each ring in isolation. With a larger channel spacing, there would be a larger range of low power transfer (and thus higher insertion loss) between each two neighboring peaks. Fig. 4 shows the insertion loss of an eight-wavelength active BOSE under temperature variations. With a quality factor Q of 5000, the thermal-induced insertion loss of an active BOSE increases significantly with the increasing channel spacing from 1 to 4.155 nm. For channel spacing of 1 nm, the worst-case insertion loss is around 7.8 dB. It increases to 13.3 dB if the channel spacing is 2.355 nm. For channel spacing of 4.155 nm, the worst-case insertion loss is around 18.2 dB when the temperature variation is 35.2 °C. Thermal-based or electronic-based adjustment should be applied to active BOSEs in order to compensate for temperature variations.

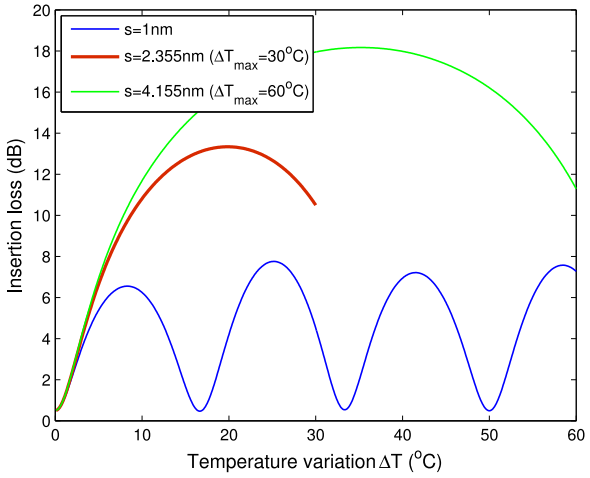

Fig. 4. Insertion loss of an eight-wavelength active BOSE with temperature variations.

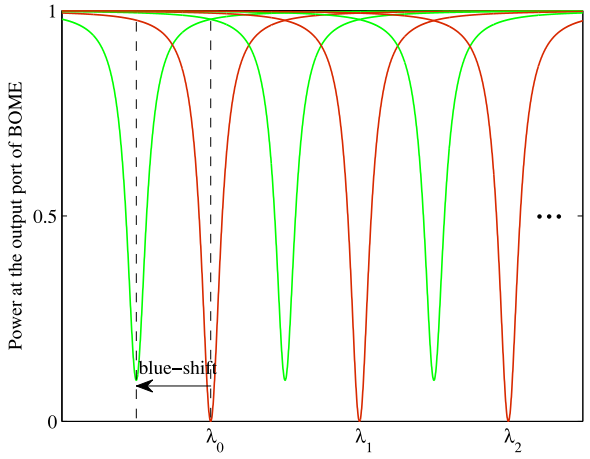

Fig. 5. Output port power transfer function of the M-wavelength BOME.

C. Thermal Model of the BOME

The BOME [Fig. 1(c)] consists of an array of M MR-based modulators which are coupled to a single waveguide. The M modulators in a off-state BOME are set to resonate at exactly the same wavelengths with the corresponding WDM laser wavelengths $\lambda_0, \lambda_1, \ldots, \lambda_{M-1}$, with a channel spacing of s between neighboring modulators (the red curves in Fig. 5). Passing laser signals to the off-state BOME would be modulated to data 0 (the output power is 0). When a modulator in BOME is turned on by forward-bias voltage, the resonance of the modulator is blue-shifted away from the laser wavelength and the passing laser signals with the corresponding wavelength would be modulated to data 1 (most power of the passing laser signals is transferred to the output port of the modulator). The power transfer at the output port of one on-state modulator is modeled as $T_{\text{modulator}}$ in (12), where δ is half of the 3 dB bandwidth of the modulators, b is the blue-shift distance when the modulator is turned on, κ^2 is the fraction of power coupling between the waveguide and the ring, and κ_p^2 is the power loss per round-trip of the ring [38]

$$T_{\text{modulator}} = \frac{b^2 + \delta^2 \cdot \left(\frac{\kappa^2 - \kappa_p^2}{\kappa^2 + \kappa_p^2}\right)^2}{b^2 + \delta^2}. \tag{12}$$

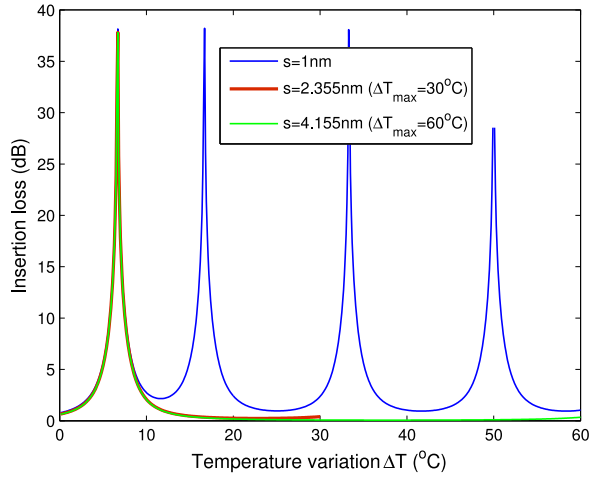

Fig. 6. Insertion loss of an eight-wavelength BOME with temperature variations.

Since the modulators in BOME is wavelength-selective, the insertion loss of BOME is temperature sensitive as a result of the thermo-optic effect. If the temperature increases from room temperature T_0 by ΔT, the resonance wavelengths of the modulators in BOME will red-shift by $\rho_{\text{MR}} \cdot \Delta T$. For optical signal with wavelength of λ_0, the worst-case insertion loss of an M-wavelength BOME is as (13), where the worst-case occurs when all the M modulators are in the on state

$$L_{\text{BOME\_0}} = \sum_{i=0}^{M-1} 10 \log \frac{\left(\frac{i \cdot s - b + \rho_{\text{MR}} \cdot \Delta T}{\delta}\right)^2 + 1}{\left(\frac{i \cdot s - b + \rho_{\text{MR}} \cdot \Delta T}{\delta}\right)^2 + \left(\frac{\kappa^2 - \kappa_p^2}{\kappa^2 + \kappa_p^2}\right)^2}. \tag{13}$$

For optical signals with wavelength $\lambda_x, x \in (0, M-1]$, the worst-case insertion loss of the M-wavelength BOME is as (14), where the modulators for wavelengths $\lambda_x, \lambda_{x+1}, \ldots \lambda_{M-1}$ are in the on state and the modulators for wavelength $\lambda_0, \lambda_1, \ldots \lambda_{x-1}$ are in the off state

$$L_{\text{BOME\_x}} = \sum_{i=0}^{M-x-1} 10 \log \frac{\left(\frac{i \cdot s - b + \rho_{\text{MR}} \cdot \Delta T}{\delta}\right)^2 + 1}{\left(\frac{i \cdot s - b + \rho_{\text{MR}} \cdot \Delta T}{\delta}\right)^2 + \left(\frac{\kappa^2 - \kappa_p^2}{\kappa^2 + \kappa_p^2}\right)^2} + \sum_{j=1}^{x} 10 \log \frac{\left(\frac{j \cdot s - \rho_{\text{MR}} \cdot \Delta T}{\delta}\right)^2 + 1}{\left(\frac{j \cdot s - \rho_{\text{MR}} \cdot \Delta T}{\delta}\right)^2 + \left(\frac{\kappa^2 - \kappa_p^2}{\kappa^2 + \kappa_p^2}\right)^2}. \tag{14}$$

Fig. 6 shows the insertion loss of an eight-wavelength BOME to the laser wavelength λ_7 (we assume the wavelength λ_7 is at 1550 nm). We assume the blue-shift distance b of the on-state BOME is 0.4 nm, and the qualify factor Q of modulators is 5000. It is shown that by using a large channel spacing of 2.355 or 4.155 nm, the thermal-aware insertion loss of the BOME can be well controlled under 3 dB except for the temperature variation range from 4 °C to 9 °C. The peak of insertion loss occurs at ΔT of 6.7 °C, where the modulator for wavelength λ_7 wrongly shifts to the off-state due to temperature increase. Thermal-based or electronic-based adjustment should be applied to BOME in order to compensate for temperature variations.

D. Thermal Model of the BOFE

The BOFE [Fig. 1(d)] is used at the receiver to demultiplex the WDM wavelengths before the optical signals reaching the photodetectors. An M-wavelength BOFE consists of M MR-based add-drop filters, where all the rings coupled to a single input waveguide. The i_{th} ($i \in [0, M-1]$) add-drop filter in the BOFE is used to filter out the laser wavelength λ_i correspondingly. Since the MR-based add-drop filter is wavelength-selective, the insertion loss of the BOFE is temperature dependent. If the working temperature of BOFE increases from room temperature T_0 by ΔT, the resonance wavelengths of the M add-drop filters in the BOFE will redshift by $\rho_{\text{MR}} \cdot \Delta T$. For optical signals with wavelength of λ_0, the insertion loss of the BOFE (15) is the drop-port insertion loss of the add-drop filter for wavelength λ_0, where δ is half of the 3 dB bandwidth of the drop-port power transfer spectrum of the add-drop filter, κ^2 is the fraction of power coupling between the input/output waveguides and the ring, and κ_p^2 is the power loss per round-trip of the rings [38]

$$L_{\text{BOFE\_0}} = 10\log\left(\left(\frac{2\kappa^2+\kappa_p^2}{2\kappa^2}\right)^2 \cdot \left(\frac{(\rho_{\text{MR}}\Delta T)^2+\delta^2}{\delta^2}\right)\right). \tag{15}$$

For optical signals with wavelength $\lambda_x, x \in [1, M-1]$, the insertion loss of the M-wavelength BOFE is the summation of the through-port insertion loss of the add-drop filters for wavelength $\lambda_0, \lambda_1, \ldots, \lambda_{x-1}$, and the drop-port insertion loss of the add-drop filter for wavelength λ_x

$$\begin{aligned} L_{\text{BOFE\_}x} &= 10\log\left(\left(\frac{2\kappa^2+\kappa_p^2}{2\kappa^2}\right)^2 \cdot \left(\frac{(\rho_{\text{MR}}\Delta T)^2+\delta^2}{\delta^2}\right)\right) \\ &+ \sum_{i=0}^{x-1} 10\log\frac{((x-i)\cdot s+\rho_{\text{MR}}\Delta T)^2+\delta^2}{((x-i)\cdot s+\rho_{\text{MR}}\Delta T)^2+\delta^2 \cdot \left(\frac{\kappa_p^2}{2\kappa^2+\kappa_p^2}\right)^2}. \end{aligned} \tag{16}$$

Fig. 7 shows the insertion loss of an eight-wavelength BOFE to laser signals with wavelength λ_7. We assume the quality factor Q of the microresonators is 5000. It is shown that even if with a large channel spacing, the thermal-induced insertion loss of the BOFE can be as high as 20 dB if the temperature variation is over 30 °C. Thermal-based or electronic-based adjustment should be applied to BOFE in order to compensate for temperature variations.

E. WDM-Based ONoC Thermal Model

In the previous sections, we presented the thermal models of each optical component in an M-wavelength WDM-based

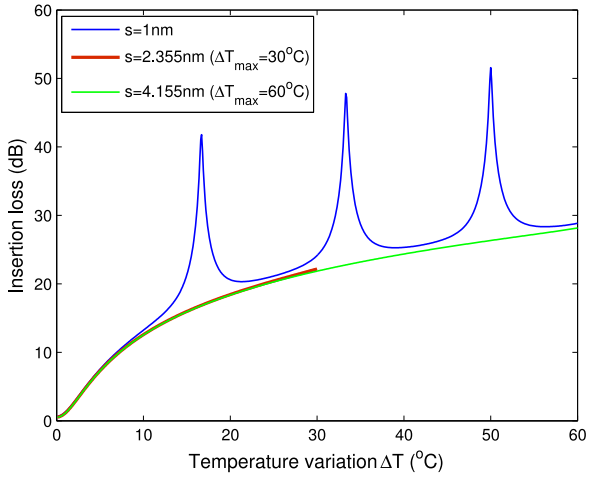

Fig. 7. Insertion loss of an eight-wavelength BOFE with temperature variations.

optical link. To ensure that ONoCs function properly, a necessary condition is that the optical signal power received by the receiver of an optical link should not be lower than the receiver sensitivity. We assume that the sensitivity is -14.2 dBm for a), bit error rate (BER) of 10^{-12} [39]. This condition must hold, otherwise the BER would increase. We model the condition in (18), where P_{TX} is the output power of the optical transmitter on the link, L_{total} is the total optical power loss in the path, and S_{RX} is the sensitivity of the receiver. For any optical transmission, the power of an optical signal generated from the laser source can be measured by adding the total power loss along the path to the minimum optical power required at the optical receiver. The total optical power loss in a link is the summation of the insertion loss of all the passing optical components, and the waveguide propagation loss. The input power at the transmitter should increase to compensate for the additional optical power loss caused by thermal variations. Thus the insertion loss of optical components directly affects the power consumption of the optical NoC

$$P_{\text{TX}} - L_{\text{total}} \geq S_{\text{RX}}. \tag{18}$$

We can get the ONoC thermal model in (17), shown at the bottom of the page. If using off-chip VCSELs as the laser sources, T_{VCSEL} can be fixed by equipping the laser sources with temperature control unit. If using on-chip VCSELs as the laser sources, T_{VCSEL} varies in the on-chip temperature range $[T_{\min}, T_{\max}]$. Under a high power loss in the optical path caused by chip temperature fluctuations, more input power would be needed by the transmitter to guarantee enough optical power reaching the receiver. We assume that the number of active switching and passive switching is N_1 and N_2

$$\begin{aligned} &10\log\left(\left(I-\alpha-\beta\left(T_{\text{VCSEL}}-T_{\text{th}}\right)^2\right)\left(\varepsilon-\gamma\cdot T_{\text{VCSEL}}\right)\right) - L_{\text{BOME\_}x} - \sum_{i=0}^{N_1-1} L_{\text{BOSE\_active\_}k} \\ &- \sum_{j=0}^{N_2-1} L_{\text{BOSE\_parking\_}j} - L_{\text{BOFE\_}x} - L_{\text{WG}} \geq S_{\text{RX}}. \end{aligned} \tag{17}$$

respectively. The total optical power loss to optical signals at wavelength $\lambda_x, x \in [0, M-1]$ is the summation of insertion loss of each optical component in the optical link, where $L_{\text{BOME\_}x}$ is the insertion loss in the modulation stage (14), $L_{\text{BOSE\_active\_}i}$ is the insertion loss in the i_{th} active switching stage (8), $L_{\text{BOSE\_parking\_}j}$ is the insertion loss in the j_{th} parking switching stage, L_{WG} is the waveguide propagation loss in the path, and $L_{\text{BOFE\_}x}$ is the insertion loss of BOFE to optical signals at wavelength λ_x (16). The thermal-aware optical power received at the end of the link can be calculated by the left side of (17).

F. Thermal-Based Adjustment

MR-based BOSE, as well as BOME and BOFE, are wavelength selective. As a result of the thermo-optic effect, microresonators red-shift with increasing temperatures. If using off-chip laser sources with temperature control, the laser wavelengths could be fixed. If using on-chip laser sources, the laser wavelengths will also shift away from the original laser wavelengths due to on-chip temperature variations. As shown from Figs. 3–7, the temperature-dependent wavelength shift results in more insertion loss in BOSEs, as well as in BOME and BOFE. By setting a larger channel spacing (e.g., 2.355 nm for 30 °C of temperature variation), the insertion loss of a parking BOSE could be well controlled under 3 dB. However, even with the large channel spacing, active BOSEs as well as BOME and BOFE would still introduce a significant insertion loss under temperature variations. Thermal-based or electronic-based adjustment can be used for active BOSEs as well as BOME and BOFE to compensate for temperature variations. For parking BOSEs, if without large channel spacing, thermal-based or electronic-based adjustment are needed to move the parking BOSEs out of the misplace region. By using electronic-based adjustment, the resonance wavelength of microresonators can be blue-shifted by at most 1 nm. Thermal-based adjustment can red-shift the resonance wavelength of microresonators, and the tuning range is on the order of tens of nm, which is much wider than the electronic-based adjustment. The tuning efficiency of thermal-based adjustment in current technology is in the order of several mW/nm [27], [28].

Fig. 8 shows an example of an active BOSE under temperature variation ΔT, where the resonance of the rings in the active BOSE shift away from the corresponding laser wavelengths by $\rho_{\text{MR}} \cdot \Delta T$ to $\lambda_{\text{MR\_}i\text{\_thermal}}$. By using the technique of thermal-based adjustment with channel remapping, the resonance of each ring of the active BOSE can be adjusted to the neighboring next laser wavelength by thermal tuning (e.g., for $\rho_{\text{MR}} \cdot \Delta T < s$, the i_{th} ring resonance is remapped to $(i+1)_{\text{th}}$). Additional guard rings are necessary for thermal-based adjustment with channel remapping. For example, if $\rho_{\text{MR}} \cdot \Delta T < s$, one guard ring would be needed to be mapped to the laser wavelength λ_0. For thermal-based adjustment with channel remapping, the tuning distance d for each ring is as

$$d = \left\lceil \frac{\rho_{\text{MR}} \cdot \Delta T}{s} \right\rceil \cdot s - \rho_{\text{MR}} \cdot \Delta T. \tag{19}$$

For the baseline setting in Fig. 2, the thermal-based adjustment must be used together with the additional guard rings,

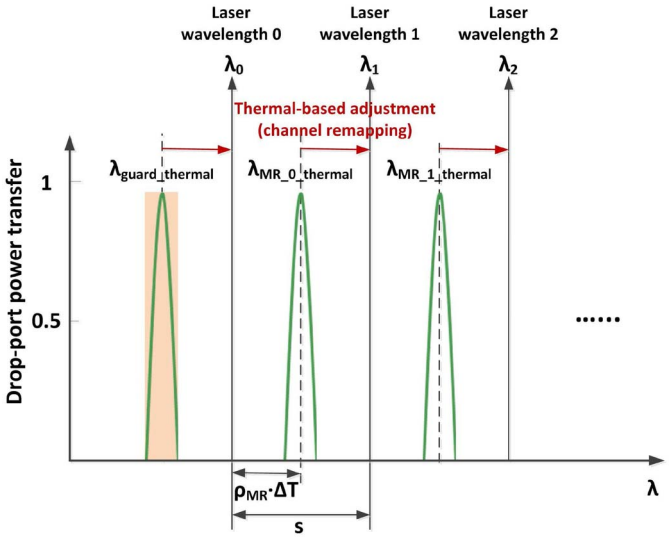

Fig. 8. Thermal-based adjustment with channel remapping.

otherwise optical signals of laser wavelength λ_0 or more would be lost. Multiple guard rings would be needed if $\rho_{\text{MR}} \cdot \Delta T > s$. For example, if $s < \rho_{\text{MR}} \cdot \Delta T < 2$ s, the i_{th} ring resonance is remapped to $(i+2)_{\text{th}}$, and two guard rings are needed for laser wavelength λ_0 and λ_1. With the guard rings and the technique of channel remapping, the tuning distance of thermal-based adjustment for each ring is well within one channel spacing. However, the additional guard rings cost more.

In order to use thermal-based adjustment without guard rings, since the thermal tuning can only do red-shift of wavelength, the resonance wavelength of a MR should always be no larger than the corresponding laser wavelength, even under the maximum chip temperature. For a maximum on-chip temperature variation of ΔT_{max}, the resonance wavelength of the MR at the maximum possible temperature should be no larger than the corresponding laser wavelength at room temperature T_0. This can be guaranteed by setting $\lambda_{\text{MR\_0}}$ as the condition shown in (20), where $\lambda_{\text{MR\_0}}$ is the resonance wavelength of the 0_{th} ring in the M-wavelength BOSE at room temperature T_0, and λ_0 is the 0_{th} laser wavelength at room temperature

$$\lambda_{\text{MR\_0}} = \lambda_0 - \rho_{\text{MR}} \cdot \Delta T_{\text{max}}. \tag{20}$$

In this case, even under the maximum temperature variation, each MR can be adjusted to the corresponding laser wavelength by thermal-based adjustment. The tuning distance is as

$$d = \rho_{\text{MR}} \cdot (\Delta T_{\text{max}} - \Delta T). \tag{21}$$

IV. OTemp: Optical Thermal Effect Modeling Platform

We developed OTemp (Fig. 9), an optical thermal effect modeling platform for both WDM-based and single-wavelength optical links in ONoCs. OTemp is based on the system-level ONoC thermal models presented in the previous sections. OTemp is a C++ based program which analyzes the thermal-aware power consumption as well as the optical power loss for optical links under temperature variations. The inputs

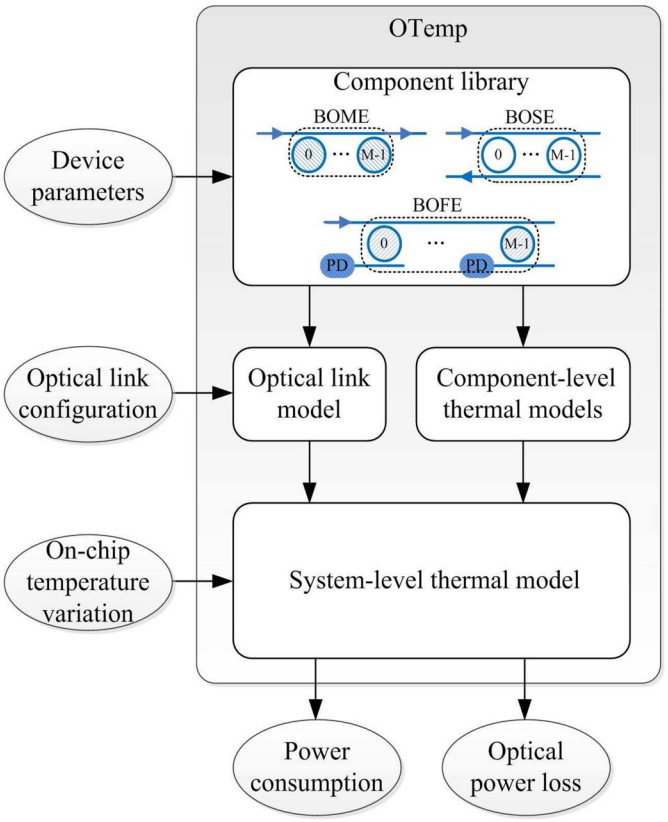

Fig. 9. OTemp: an optical thermal effect modeling platform.

to this tool include the on-chip temperature range $[T_{\min}, T_{\max}]$, the optical link configurations and photonic devices parameters. The optical link configurations include the number of wavelengths in WDM; WDM channel spacing; off-chip or on-chip laser; direct modulation or using BOME; modulation speed; the number of active BOSEs; the number of parking BOSEs; with or without channel remapping for thermal-based adjustment; use the by-default or the optimal initial setting of MRs; electronic-based or thermal-based switching; the power consumption of the laser driver, serializer, deserializer, TIA-LA; the power consumption of the thermal-based adjustment; the receiver sensitivity; the waveguide propagation loss; the waveguide crossing loss, etc. OTemp models the thermal effects in system-level as well as in component level. The component library includes optical link components such as the BOSE, BOME, and BOFE.

A. Analysis of WDM-Based Optical Links

The configurations for a WDM-based optical link includes choosing whether to use on-chip or off-chip VCSEL as the laser source (by-default: use off-chip VCSEL); choosing whether to use BOME for modulation or use direct-modulated VCSEL (by-default: use BOME for modulation); choosing whether to use guard rings for channel remapping in thermal-based adjustment (by-default: with guard rings and channel remapping); choosing whether to use the by-default setting or the optimal setting of λ_{MR_0} (by-default: $\lambda_{MR\_0} = \lambda_0$); the WDM channel spacing (by-default: 1 nm), the number of WDM wavelength (by-default: 8); the switching mechanism for BOSE (by-default: electronic-based switching), the number of active BOSE in the link (by-default: 3); the number of parking BOSE in the link (by-default: 10); and the quality factor of the microresonators used in optical components (by-default: 5000).

For analysis of WDM-based optical links, the outputs include the worst-case thermal-aware power consumption with or without thermal-based adjustments. For the case with thermal-based adjustment, two scenarios are analyzed including using the thermal-based adjustment together with guard rings for channel remapping, or using the thermal-based adjustment without channel remapping. The by-default setting of λ_{MR_0} (resonance wavelength of the 0_{th} ring in the M-wavelength BOSE and BOFE and BOME at room temperature T_0) is equal to λ_0 (the 0_{th} WDM laser wavelength at room temperature). If using thermal-based adjustment without channel remapping, since the thermal tuning can only do red-shift of wavelength, the resonance wavelength of a MR should always be no larger than the corresponding laser wavelength, even under the maximum chip temperature. This can be guaranteed by setting λ_{MR_0} as (20), according to which λ_{MR_0} is lower than λ_0 and the difference is determined by the maximum temperature variation $\Delta T_{\max}$. The output results include both of the total power consumption and the on-chip power consumption. If using on-chip laser as the laser source, the total power consumption is equal to the on-chip power consumption. If using off-chip laser as the laser source, the on-chip power consumption does not include the power consumed by the laser source itself.

B. Analysis of Single-Wavelength Optical Links

For analysis of single-wavelength optical links, the outputs include the worst-case and average-case thermal-aware power consumption under different combinations of low-temperature-dependence techniques, such as the optimal setting of λ_{MR_0} (22), thermal-based adjustment for the active switching stages, and the use of athermal microresonators. The worst-case and average-case analysis are conducted among all possible thermal maps where the chip temperature varies spatially between $T_{\min}$ and $T_{\max}$. If using on-chip VCSELs as the laser source, the worst-case occurs when the VCSEL works at the maximum on-chip temperature $T_{\max}$ while microresonators at all the active switching stages work at the minimum on-chip temperature $T_{\min}$. In this case, the laser power efficiency would be the worst, and the wavelength mismatch between the laser and the active switching stages would also be the worst. The average-case power consumption is calculated as the mathematical expectation of the power consumption, while assuming that the temperatures of on-chip VCSEL and all microresonators are all uniformly distributed between $T_{\min}$ and $T_{\max}$. The by-default setting of MR wavelength at room temperature is $\lambda_{MR\_0} = \lambda_{VCSEL\_0}$. The additional optical power loss under temperature variation can be reduced if we set λ_{MR_0} according to the optimal setting shown in (22) [34]

$$\lambda_{MR\_0} = \lambda_{VCSEL\_0} + \frac{(\rho_{VCSEL} - \rho_{MR})}{2} \cdot (T_{\max} + T_{\min} - 2T_0). \tag{22}$$

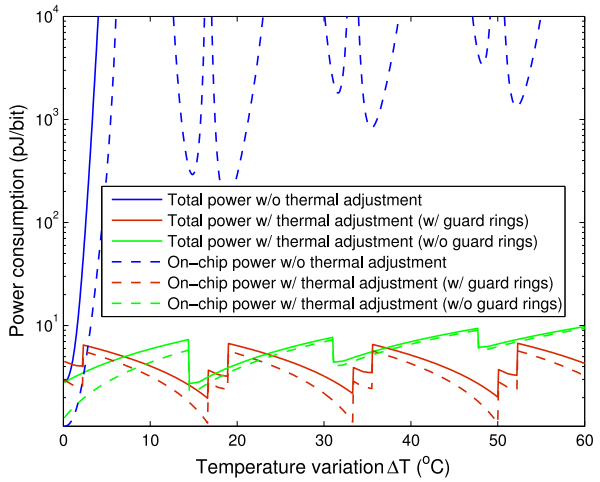

Fig. 10. Worst-case power consumption for one wavelength in an eight-wavelength WDM optical link, with off-chip laser source, and channel spacing 1 nm.

V. Case Studies and Analysis

Based on the proposed optical thermal effect modeling platform OTemp, we simulated the worst-case power consumption for an eight-wavelength WDM-based optical link under temperature variations. The case studies are all based on the abstracted optical link model (shown in Fig. 1) with different assumptions (e.g., laser source, channel spacing, with/without thermal adjustment, etc.). We assume the use of VCSELs as the laser source. The modeling of power consumption of the laser source is based on the light-current-voltage (LIV) characteristics of VCSELs demonstrated in [9]. For the O/E interface (optical-to-electronic conversion, and verse visa), we assume the use of the serializer and deserializer designs in [40], and the VCSEL driver and TIA-LA designs in [41]. We assume the electronic part of the ONoC is in 45 nm, and scale all the related power consumption linearly to 45 nm. For example, the driver and TIA-LA circuits power consumption is scaled from 0.2 pJ/bit and 0.6 pJ/bit in 80 nm to 0.1125 pJ/bit and 0.3375 pJ/bit in 45 nm, and the power consumption of serializer and deserializer is scaled from 0.576 pJ/bit in 90 nm to 0.288 pJ/bit in 45 nm. The photodetector model is based on a Ge waveguide photodetector monolithically integrated in 130 nm CMOS process with a sensitivity of -14.2 dBm for 10^{-12} of BER [39]. For the thermal-based adjustment, we assume the tuning efficiency is 3.5 mW/nm [27], [28].

A. With Off-Chip Laser Sources

Fig. 10 shows the worst-case total power consumption and on-chip power consumption for wavelength λ_7 which is (1550 nm) in an eight-wavelength WDM-based optical link with WDM channel spacing of 1 nm. We also tried to plot the figure using a much larger scale for the y-axis, which can show more data of the solid blue curve and the blue dotted curve. However, we find that the current figure show more clearly the comparison of the worst-case power consumption with different configurations. The proposed thermal model is general enough that can be applied to any thermal maps. The major contribution of this paper is the proposed thermal model, which is general enough that can be applied to any thermal maps. The purpose of showing the simulation results in this section is to demonstrate the proposed model, and demonstrate how the power consumption is affected by thermal variations. In this paper, we conducted the thermal-aware power analysis for optical NoCs for a maximum temperature variation up to 60 °C. We believe that results under a wide temperature range would give more information. We assume the laser wavelengths are fixed by using off-chip VCSEL laser source equipped with temperature control (the total power consumption does not include the power consumed by the temperature control unit). BOME is used for modulation (10 Gb/s for each wavelength). The initial setting of the BOSEs is according to the baseline setting shown in Fig. 2, with channel spacing of 1 nm. The switching of the BOSEs are electronic-based, with $\Delta\lambda_{\text{off}-\text{on}}$ of 0.4 nm. The width of the misplace region $\Delta\lambda_{\text{mis}}$ is set as three times of the MR 3 dB bandwidth 2δ, which corresponds to around 0.46 dB optical insertion loss of a parking BOSE when it is placed at the edge of the misplace region of laser wavelengths. We assume the quality factor Q of the microresonators is 5000. We assume the number of active BOSEs in the link is set as three, and the number of parking BOSEs is set as 10. For the case where no thermal-based adjustment is used, we assume an increase in the output optical power of the laser sources to compensate for the extra optical power loss inserted by the optical components under temperature variations. Since we assume a high quality factor for the microresonators in BOSE and BOME and BOFE, the worst-case power consumption increases dramatically with the temperature variation ΔT. For the case using thermal-based adjustment with guard rings, the rings resonance are remapped to the neighboring next laser wavelengths by thermal tuning. For the case using thermal-based adjustment without guard rings, the initial setting of the microresonators is according to the optimal condition shown in (20), which guarantees that the thermal-based adjustment is always applicable to tune the rings into the corresponding laser wavelengths. The power consumption is determined by the total optical power loss in the link, which is the summation of the insertion loss of all the active BOSEs, passive BOSEs, BOME, BOFE in the link. Since the insertion loss of each optical component is nonmonotonic with increasing temperatures, the power consumption of the link is nonmonotonic.

Fig. 10 shows the worst-case power consumption of wavelength λ_7 in an eight-wavelength optical link. We assume the wavelength λ_7 is at 1550 nm, and the wavelength $\lambda_i (0 \leq i < 7)$ is $1550 - i \cdot$ channel_spacing. If using thermal-based adjustment with channel remapping, the worst-case on-chip and total power consumption for wavelength λ_7 is 5.7 pJ/bit and 6.7 pJ/bit, respectively. The eight wavelengths do not have exactly the same worst-case power consumption. According to our results, the worst-case power consumption for the eight wavelengths differ slightly, but follow the same trend. For the worst-case total power consumption, the average of the eight wavelengths is 6.4 pJ/bit, which is slightly smaller than the worst-case power consumption of the wavelength λ_7. For the worst-case on-chip power consumption, the average of the eight wavelengths is 5.6 pJ/bit, which is almost the

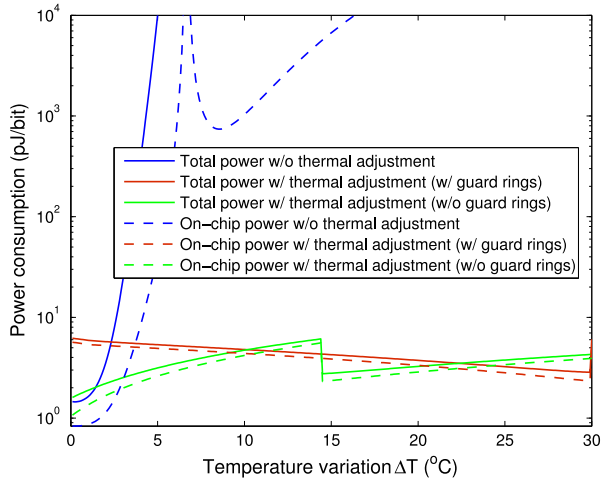

Fig. 11. Worst-case power consumption for one wavelength in an eight-wavelength WDM optical link, with off-chip laser source, and channel spacing 2.665 nm.

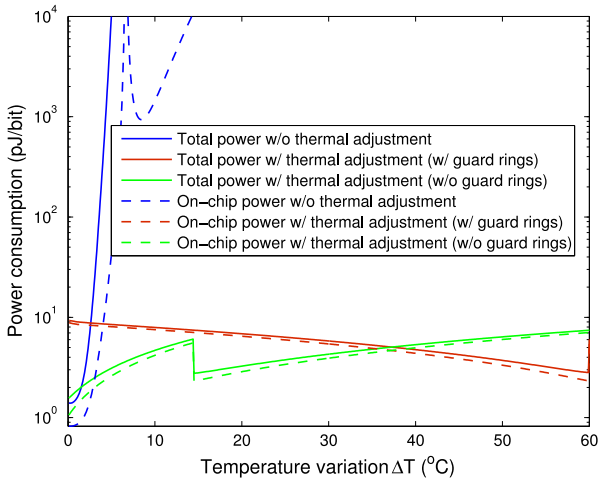

Fig. 12. Worst-case power consumption for one wavelength in an eight-wavelength WDM optical link, with off-chip laser source, and channel spacing 4.465 nm.

same with the worst-case power consumption of the wavelength λ_7. Different from the worst-case, the average-case power consumption can be evaluated while assuming that the on-chip temperature variation at any spot is randomly distributed according to a uniform distribution between 0 and 60 °C. The average-case total power consumption is around 3.6 pJ/bit, and the average-case on-chip power consumption is around 3.2 pJ/bit. The average-case power consumption is about 56% of the worst case.

If using thermal-based adjustment without guard rings but together with the optimal initial setting (20), Fig. 10 shows that the worst-case on-chip and total power consumption is respectively 9.4 and 9.8 pJ/bit for a temperature variation of 60 °C. It includes around 7 pJ/bit for the thermal based adjustment for a temperature variation of 60 °C. The laser power consumption is around 2 pJ/bit, and it costs 0.8 pJ/bit for the driver, serializer/deserializer, photodetector and TIA-LA circuits. Our results are based on the recent demonstrated values of the related devices in current technology. It has the potential for a better power efficiency as the technology of optical interconnect continues to make progress.

Fig. 11 shows the worst-case total power consumption and on-chip power consumption for wavelength λ_7 in an eight-wavelength WDM-based optical link with channel spacing of 2.665 nm. The channel spacing is chosen as 2.665 nm to guarantee that, the worst-case insertion loss of parking BOSEs is no larger than 0.46 dB for the maximum temperature variation of 30 °C. Thus Fig. 11 is plotted only until to a temperature variation of 30 °C. Results show that the worst-case on-chip and total power consumption is 5.7 and 6.2 pJ/bit respectively if using thermal-based adjustment with channel remapping. If using thermal-based adjustment without channel remapping, the worst-case on-chip and total power consumption is respectively 5.6 and 6.1 pJ/bit if following the optimal setting shown in (20).

Fig. 12 shows the worst-case total power consumption and on-chip power consumption for wavelength λ_7 in an eight-wavelength WDM-based optical link with channel spacing of 4.465 nm. The channel spacing is chosen as 4.465 nm

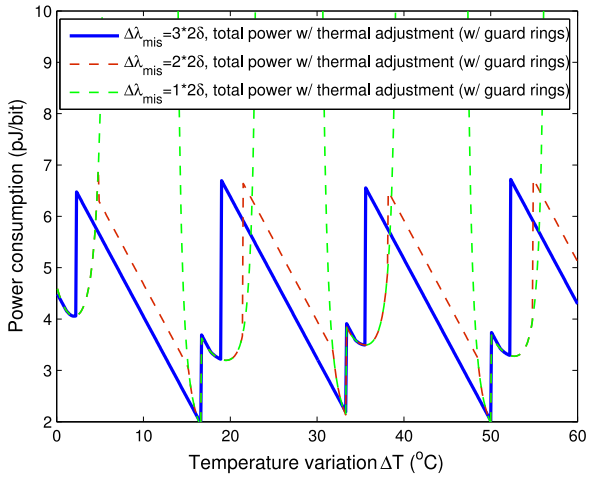

Fig. 13. Worst-case power consumption for one wavelength in an eight-wavelength WDM optical link, with different settings of $\Delta\lambda_{mis}$, w/ off-chip laser source, and channel spacing 1 nm.

to guarantee that, for the maximum temperature variation of 60 °C, the worst-case insertion loss of parking BOSEs is no larger than 0.46 dB. Results show that the worst-case on-chip and total power consumption is 8.8 and 9.3 pJ/bit respectively if using thermal-based adjustment with channel remapping. If using thermal-based adjustment without guard rings but following the optimal setting (20), the worst-case on-chip and total power consumption is respectively 7.1 and 7.4 pJ/bit for a temperature variation of 60 °C. As compared to the case with channel spacing of 1 nm (Fig. 10), we can observe that if we use thermal-based adjustment with channel remapping, a larger channel spacing of 4.465 nm would result in a larger worst-case power consumption. However, if we use thermal-based adjustment without channel remapping, the worst-case power consumption benefits from the larger channel spacing.

Parking BOSEs under temperature variations may wrongly shift into the neighboring laser wavelengths. Parking BOSEs in the misplace region would result in significant insertion loss or even switch the optical signals to a wrong direction. Thermal-based adjustment can be used to move the parking

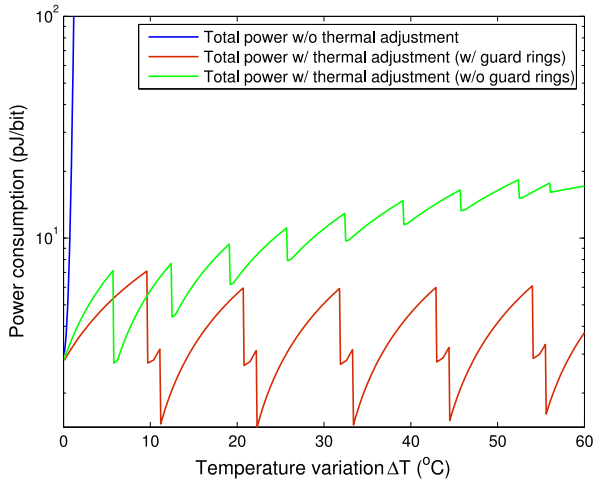

Fig. 14. Worst-case total power consumption for one wavelength in an eight-wavelength WDM optical link, with on-chip laser source, and channel spacing 1 nm.

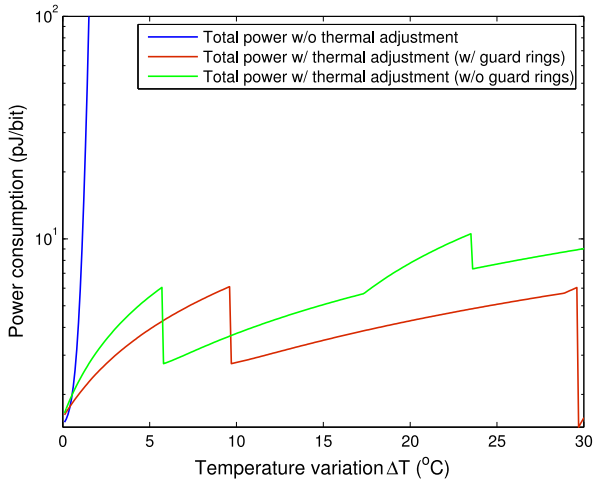

Fig. 15. Worst-case total power consumption for one wavelength in an eight-wavelength WDM optical link, with on-chip laser source, and channel spacing 2.665 nm.

BOSEs out of the misplace region. We defined the width of the misplace region of each laser wavelength as $\Delta\lambda_{\text{mis}}$ (Fig. 2). By setting a large $\Delta\lambda_{\text{mis}}$ for thermal-based adjustment, the insertion loss of parking BOSEs can be well controlled even under large temperature variations. In Figs. 10–12, the width of the misplace region $\Delta\lambda_{\text{mis}}$ is set as three times that of 3 dB bandwidth 2δ, which corresponds to around 0.46 dB optical power loss of a parking BOSE when its resonance wavelengths are at the edge of the misplace region of the laser wavelengths. Fig. 13 shows the worst-case power consumption with different settings of $\Delta\lambda_{\text{mis}}$ from three times of 2δ, two times of 2δ, to 2δ. We assume thermal-based adjustment is used together with guard rings for channel remapping. When the $\Delta\lambda_{\text{mis}}$ is set as wide as the 3 dB bandwidth of the microresonators, the insertion loss is around 3 dB when the resonance of the parking BOSE is at the edge of the misplace region. Since we assume a large number of parking BOSEs in an optical link, it will result in a significant total optical power loss if a single parking BOSE introduces 3 dB optical power loss. Fig. 13 shows that the worst-case power consumption with $\Delta\lambda_{\text{mis}}$ of $3 * 2\delta$ or $2 * 2\delta$ is reduced significantly than the power consumption with $\Delta\lambda_{\text{mis}}$ of 2δ.

B. With On-Chip Laser Source

Figs. 14–16 show the worst-case total power consumption for wavelength λ_7 in an eig-wavelength WDM-based optical link with on-chip VCSELs as the WDM laser sources. We assume the wavelength λ_7 is at 1550 nm, and the wavelength $\lambda_i (0 \leq i < 7)$ is $1550 - i \cdot$ channel_spacing. Compared to the case with off-chip laser source, the use of on-chip laser source would cause more on-chip power consumption. We take the degradation of VCSEL power efficiency under high temperatures into account when analyzing the thermal-aware power consumption. The thermal sensitivity of VCSEL is based on the LIV characteristics of VCSELs demonstrated in [9]. If using thermal-based adjustment with guard rings for channel remapping, with channel spacing of 1 nm, the worst-case total power consumption for wavelength λ_7 is 7.1 pJ/bit. The eight wavelengths do not have exactly the same

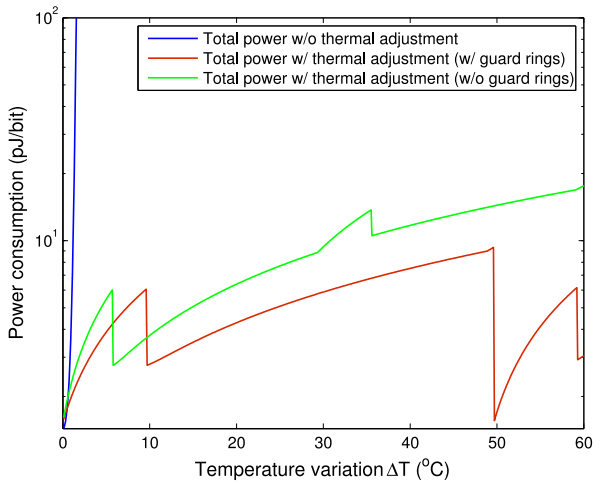

Fig. 16. Worst-case total power consumption for one wavelength in an eight-wavelength WDM optical link, with on-chip laser source, and channel spacing 4.465 nm.

worst-case power consumption, but follow the same trend. For the worst-case total power consumption, the average of the eight wavelengths is 6.9 pJ/bit, which is almost the same with the worst-case power consumption of the wavelength λ_7. The average-case total power consumption is evaluated while assuming that the on-chip temperature variation at any spot is randomly distributed according to a uniform distribution between 0 and 60 °C. The average-case power consumption is 3.5 pJ/bit, which is around half of the worst-case power consumption. If using thermal-based adjustment without guard rings but following the optimal setting (20), Fig. 14 shows that the worst-case total power consumption is 17.2 pJ/bit for a temperature variation of 60 °C.

Fig. 15 shows that with channel spacing of 2.665 nm, the worst-case total power consumption is 6.1 pJ/bit by using thermal-based adjustment with channel remapping. If using thermal-based adjustment without guard rings, the worst-case total power consumption is 10.5 pJ/bit for a temperature variation of 30 °C when following the optimal setting (20). Fig. 16 shows that with channel spacing of 4.465 nm, the

worst-case total power consumption is 9.1 pJ/bit if using thermal-based adjustment with channel remapping. If using thermal adjustment without guard rings, the worst-case total power consumption is 17.2 pJ/bit for a temperature variation of 60 °C. Compared to the results with channel spacing of 1 nm (Fig. 14), it shows that if using thermal-based adjustment with guard rings for channel remapping, larger channel spacing would result in a larger worst-case power consumption. If using thermal-based adjustment without channel remapping but following the optimal setting (20), the worst-case power consumption is similar when using larger channel spacing.

VI. Conclusion

For the first time, this paper systematically modeled and analyzed the thermal effects in WDM-based ONoCs; and developed OTemp simulator, an optical thermal effect modeling platform for both WDM-based and single-wavelength optical links. We use case studies to analyze the worst-case power consumption for one wavelength in an eight-wavelength WDM-based optical link with different configurations. Results show that the worst-case power consumption increases dramatically with temperature variations. Thermal-based adjustment and optimal device settings can help reduce the worst-case power consumption. It shows that the use of thermal-based adjustment can reduce the worst-case power consumption to less than 10 pJ/bit. We can observe that there is a tradeoff in using thermal-based adjustment (e.g., with guard rings for channel remapping, or without guard rings but with the optimal initial setting of microresonators). The use of thermal-based adjustment with guard rings for channel remapping has a better effect of power reduction, but requires a higher cost for guard rings. From the set of simulation results under different channel spacing, we can also learn that a larger channel spacing could not simply reduce the impacts of thermal effects in WDM-based ONoCs. For example, if we use thermal-based adjustment with channel remapping, a larger channel spacing (e.g., 4.465 nm) would result in a larger worst-case power consumption. For future research work in WDM-based ONoC designs, thermal sensitivity is a key issue. The proposed thermal model and simulation results shown in this paper can help designers fully understand the thermal effects in ONoCs, and provide a practical guide in the early design stage to reduce the thermal sensitivity by low-temperature-dependence techniques. In addition, the OTemp simulator can be used by ONoC designers to quantitatively analyze the link-level power consumption under temperature variations for their ONoC designs.

References

[1] W. Dally and B. Towles, "Route packets, not wires: On-chip interconnection networks," in *Proc. Design Autom. Conf.*, 2001, pp. 684–689.

[2] L. Benini and G. De Micheli, "Networks on chip: A new paradigm for systems on chip design," in *Proc. Design Autom. Test Eur. Conf. Exhibit.*, Paris, France, 2002, pp. 418–419.

[3] E. Bonetto, L. Chiaraviglio, D. Cuda, G. Gavilanes Castillo, and F. Neri, "Optical technologies can improve the energy efficiency of networks," in *Proc. 35th Eur. Conf. Opt. Commun.*, Vienna, Austria, 2009, pp. 1–4.

[4] I. O'Connor, "Optical solutions for system-level interconnect," in *Proc. Int. Workshop Syst. Level Interconnect Prediction*, New York, NY, USA, 2004, pp. 79–88.

[5] I. Young *et al.*, "Optical I/O technology for tera-scale computing," in *Proc. IEEE Int. Solid-State Circuits Conf.*, San Francisco, CA, USA, 2009, pp. 468–469.

[6] A. Shacham, K. Bergman, and L. Carloni, "On the design of a photonic network-on-chip," in *Proc. 1st Int. Symp. Netw. Chip*, Washington, DC, USA, 2007, pp. 53–64.

[7] H. Gu, J. Xu, and W. Zhang, "A low-power fat tree-based optical network-on-chip for multiprocessor system-on-chip," in *Proc. Design Autom. Test Eur. Conf. Exhibit.*, Nice, France, Apr. 2009, pp. 3–8.

[8] Y. Pan *et al.*, "Firefly: Illuminating future network-on-chip with nanophotonics," in *Proc. 36th Int. Symp. Comput. Architecture (ISCA)*, Austin, TX, USA, 2009, pp. 429–440.

[9] A. Syrbu *et al.*, "10 Gbps VCSELs with high single mode output in 1310 nm and 1550 nm wavelength bands," in *Proc. Conf. Opt. Fiber Commun./Nat. Fiber Opt. Eng.*, San Diego, CA, USA, 2008, pp. 1–3.

[10] K. Skadron *et al.*, "Temperature-aware microarchitecture: Modeling and implementation," *ACM Trans. Archit. Code Optim.*, vol. 1, no. 1, pp. 94–125, 2004.

[11] F. Li *et al.*, "Design and management of 3D chip multiprocessors using network-in-memory," in *Proc. 33rd Int. Symp. Comput. Archit.*, Boston, MA, USA, 2006, pp. 130–141.

[12] S. Selvaraja, W. Bogaerts, P. Dumon, D. Van Thourhout, and R. Baets, "Subnanometer linewidth uniformity in silicon nanophotonic waveguide devices using CMOS fabrication technology," *IEEE J. Sel. Topics Quantum Electron.*, vol. 16, no. 1, pp. 316–324, Jan./Feb. 2010.

[13] K. H. Mo *et al.*, "A hierarchical hybrid optical-electronic network-on-chip," in *Proc. IEEE Comput. Soc. Annu. Symp. VLSI*, Lixouri, Greece, Jul. 2010, pp. 327–332.

[14] D. Ding, Y. Zhang, H. Huang, R. Chen, and D. Pan, "O-Router: An optical routing framework for low power on-chip silicon nanophotonic integration," in *Proc. 46th ACM/IEEE Design Autom. Conf.*, San Francisco, CA, USA, Jul. 2009, pp. 264–269.

[15] D. Vantrease *et al.*, "Corona: System implications of emerging nanophotonic technology," in *Proc. 35th Int. Symp. Comput. Archit.*, Beijing, China, 2008, pp. 153–164.

[16] N. Kirman *et al.*, "On-chip optical technology in future bus-based multicore designs," *IEEE Micro*, vol. 27, no. 1, pp. 56–66, Jan./Feb. 2007.

[17] M. Briere *et al.*, "System level assessment of an optical NoC in an MPSoC platform," in *Proc. Design Autom. Test Eur. Conf. Exhibit.*, Nice, France, 2007, pp. 1–6.

[18] R. Beausoleil *et al.*, "A nanophotonic interconnect for high-performance many-core computation," in *Proc. 16th IEEE Symp. High Perform. Interconnects*, Stanford, CA, USA, Aug. 2008, pp. 182–189.

[19] S. Pasricha and N. Dutt, "ORB: An on-chip optical ring bus communication architecture for multi-processor systems-on-chip," in *Proc. Asia South Pac. Conf. Design Autom.*, Seoul, Korea, Mar. 2008, pp. 789–794.

[20] M. J. Cianchetti, J. C. Kerekes, and D. H. Albonesi, "Phastlane: A rapid transit optical routing network," in *Proc. 36th Annu. Int. Symp. Comput. Archit.*, New York, NY, USA, 2009, pp. 441–450.

[21] A. Kodi, R. Morris, A. Louri, and X. Zhang, "On-chip photonic interconnects for scalable multi-core architectures," in *Proc. 3rd ACM/IEEE Int. Symp. Netw. Chip*, San Diego, CA, USA, May 2009, p. 90.

[22] C. Batten *et al.*, "Building manycore processor-to-DRAM networks with monolithic silicon photonics," in *Proc. 16th IEEE Symp. High Perform. Interconnects*, Stanford, CA, USA, 2008, pp. 21–30.

[23] R. K. Dokania and A. B. Apsel, "Analysis of challenges for on-chip optical interconnects," in *Proc. 19th ACM Great Lakes Symp. VLSI*, New York, NY, USA, 2009, pp. 275–280.

[24] T. Baehr-Jones *et al.*, "Analysis of the tuning sensitivity of silicon-on-insulator optical ring resonators," *J. Lightw. Technol.*, vol. 23, no. 12, pp. 4215–4221, Dec. 2005.

[25] V. Padgaonkar and A. Arbor, "Thermal effects in silicon based resonant cavity devices," *NNIN REU Research Accomplishments*, 2004.

[26] P. Dumon *et al.*, "Linear and nonlinear nanophotonic devices based on silicon-on-insulator wire waveguides," *Jpn. J. Appl. Phys.*, vol. 45, pp. 6589–6602, Aug. 2006.

[27] F. Gan *et al.*, "Maximizing the thermo-optic tuning range of silicon photonic structures," in *Proc. Photon. Switching*, San Francisco, CA, USA, Aug. 2007, pp. 67–68.

[28] M. Geng *et al.*, "Four-channel reconfigurable optical add-drop multiplexer based on photonic wire waveguide," *Opt. Express*, vol. 17, no. 7, pp. 5502–5516, Mar. 2009.

[29] J.-M. Lee, D.-J. Kim, H. Ahn, S.-H. Park, and G. Kim, "Temperature dependence of silicon nanophotonic ring resonator with a polymeric overlayer," *J. Lightw. Technol.*, vol. 25, no. 8, pp. 2236–2243, Aug. 2007.

[30] R. Michalzik and K. J. Ebeling, "Operating principles of VCSELs," in *Vertical-Cavity Surface-Emitting Laser Devices*. Berlin, Germany: Springer, 2003.
[31] S. Mogg *et al.*, "Temperature sensitivity of the threshold current of long-wavelength InGaAs-GaAs VCSELs with large gain-cavity detuning," *IEEE J. Quantum Electron.*, vol. 40, no. 5, pp. 453–462, May 2004.
[32] S. Koester, L. Schares, C. Schow, G. Dehlinger, and R. John, "Temperature-dependent analysis of Ge-on-SOI photodetectors and receivers," in *Proc. 3rd IEEE Int. Conf. Group IV Photon.*, Ottawa, ON, Canada, 2006, pp. 179–181.
[33] M. Morse, O. Dosunmu, G. Sarid, and Y. Chetrit, "Performance of Ge-on-Si p-i-n photodetectors for standard receiver modules," *IEEE Photon. Technol. Lett.*, vol. 18, no. 23, pp. 2442–2444, Dec. 1, 2006.
[34] Y. Ye *et al.*, "System-level modeling and analysis of thermal effects in optical networks-on-chip," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 21, no. 2, pp. 292–305, Feb. 2013.
[35] F. G. Della Corte, M. Esposito Montefusco, L. Moretti, I. Rendina, and G. Cocorullo, "Temperature dependence analysis of the thermo-optic effect in silicon by single and double oscillator models," *J. Appl. Phys.*, vol. 88, no. 12, pp. 7115–7119, Dec. 2000.
[36] B. E. Little, S. T. Chu, J. V. Hryniewicz, and P. P. Absil, "Filter synthesis for periodically coupled microring resonators," *Opt. Lett.*, vol. 25, no. 5, pp. 344–346, 2000.
[37] M.-C. M. Lee and M. C. Wu, "Variable bandwidth of dynamic add-drop filters based on coupling-controlled microdisk resonators," *Opt. Lett.*, vol. 31, no. 16, pp. 2444–2446, 2006.
[38] S. Xiao, M. H. Khan, H. Shen, and M. Qi, "Modeling and measurement of losses in silicon-on-insulator resonators and bends," *Opt. Exp.*, vol. 15, no. 17, pp. 10553–10561, 2007.
[39] G. Masini, G. Capellini, J. Witzens, and C. Gunn, "A 1550nm, 10 Gbps monolithic optical receiver in 130 nm CMOS with integrated Ge waveguide photodetector," in *Proc. 4th IEEE Int. Conf. Group IV Photon.*, Tokyo, Japan, 2007, pp. 1–3.
[40] J. Poulton *et al.*, "A 14-mW 6.25-Gb/s transceiver in 90-nm CMOS," *IEEE J. Solid-State Circuits*, vol. 42, no. 12, pp. 2745–2757, Dec. 2007.
[41] C. Kromer *et al.*, "A 100-mW 4 x 10 Gb/s transceiver in 80-nm CMOS for high-density optical interconnects," *IEEE J. Solid-State Circuits*, vol. 40, no. 12, pp. 2667–2679, Dec. 2005.

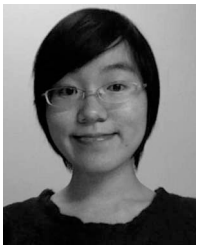

Yaoyao Ye (S'09–M'13) received the B.S. degree from the University of Science and Technology of China, Hefei, China, in 2008, and the Ph.D. degree in electronic and computer engineering from the Hong Kong University of Science and Technology, Hong Kong, in 2013.

Her current research interests include network-on-chip, inter/intrachip optical interconnect, multiprocessor system-on-chip, and embedded system.

Zhehui Wang (S'12) received the B.S. degree in electrical engineering from Fudan University, China, in 2010. He is currently pursuing the Ph.D. degree from the Department of Electronic and Computer Engineering, Hong Kong University of Science and Technology, Hong Kong.

His current research interests include embedded system, multiprocessor systems, network-on-chip, and floorplan design for network-on-chip.

Peng Yang (S'14) recieved the B.Sc. degree in electronic science and technology from Wuhan University, China, in 2013. He is currently pursuing the Ph.D. degree from the Department of Electronic and Computer Engineering, Hong Kong University of Science and Technology, Hong Kong.

His current research interests include embedded system, multiprocessor systems, and network-on-chip.

Jiang Xu (S'02–M'07) received the Ph.D. degree from Princeton University, Princeton, NJ, USA, in 2007.

He joined Bell Labs, Murray Hill, NJ, USA, as a Research Associate from 2001 to 2002. He was a Research Associate at NEC Laboratories America, Princeton, NJ, USA, from 2003 to 2005. He joined Sandbridge Technologies, USA—a startup company from 2005 to 2007, where he developed and implemented two generations of NoC-based ultralow power multiprocessor systems-on-chip for mobile platforms. He is an Associate Professor with the Hong Kong University of Science and Technology (HKUST), Hong Kong. He is the Founding Director of Xilinx-HKUST Joint Laboratory, HKUST and established Mobile Computing System Laboratory, HKUST. His current research interests include network-on-chip, multiprocessor system-on-chip, optical interconnects, embedded system, computer architecture, low-power VLSI design, and HW/SW codesign. He has authored and co-authored over 80 book chapters and papers in peer-reviewed journals and international conferences.

Dr. Xu currently serves as an Area Editor of NoC, SoC, and GPU for *ACM Transactions on Embedded Computing Systems* and an Associate Editor for the IEEE TRANSACTION ON VERY LARGE SCALE INTEGRATION SYSTEMS. He is an ACM Distinguished Speaker and an IEEE Distinguished Lecturer. He served on the steering committees, organizing committees, and technical program committees of several international conferences.

Xiaowen Wu (S'12) received the B.Sc. degree in computer science from the Harbin Institute of Technology, China, in 2008. He is currently pursuing the Ph.D. degree in electronic and computer engineering from the Hong Kong University of Science and Technology, Hong Kong.

His current research interests include embedded systems, multiprocessor systems, and network-on-chip.

Xuan Wang (S'12) received the B.S. degree in electrical engineering from Shanghai Jiaotong University, Shanghai, China, in 2009. He is currently pursuing the Ph.D. degree from the Department of Electronic and Computer Engineering, Hong Kong University of Science and Technology, Hong Kong.

His current research interests include embedded system, multiprocessor system, network-on-chip, and fault tolerant design and reliability issues in very deep submicron technologies.

Mahdi Nikdast (S'10–M'13) received the B.Sc. degree in computer engineering from Islamic Azad University, Esfahan, Iran and the Ph.D. degree from the Department of Electronic and Computer Engineering, Hong Kong University of Science and Technology, Hong Kong, in 2009 and 2013, respectively.

His current research interests include embedded and computing systems, multiprocessor systems-on-chip, networks-on-chip, and optical interconnection networks.

Zhe Wang (S'12) received the B.S. degree in electronic engineering from Shanghai Jiao Tong University, Shanghai, China, in 2011. He is currently pursuing the Ph.D. degree in electronic and computer engineering from the Hong Kong University of Science and Technology, Hong Kong.

His current research interests include embedded systems, multiprocessor system-on-chip, network-on-chip, hardware/software co-design, and design space exploration techniques.

Luan H.K. Duong (S'13) received the B.S. degree in computer science from the Hong Kong University of Science and Technology, Hong Kong, in 2012, where he is currently pursuing the Ph.D. degree in electronic and computer engineering.

His current research interests include embedded system, system-on-chip, and network-on-chip.