Hardware-software collaborative thermal sensing in optical network-on-chip-based manycore systems

Li, Mengquan; Liu, Weichen; Guan, Nan; Xie, Yiyuan; Ye, Yaoyao

2019

Li, M., Liu, W., Guan, N., Xie, Y., & Ye, Y. (2019). Hardware-software collaborative thermal sensing in optical network-on-chip-based manycore systems. *ACM Transactions on Embedded Computing Systems*, 18(6), 118:1-118:24. doi:10.1145/3362099

<https://hdl.handle.net/10356/145295>

<https://doi.org/10.1145/3362099>

© 2019 Association for Computing Machinery (ACM). All rights reserved. This paper was published in *ACM Transactions on Embedded Computing Systems* and is made available with permission of Association for Computing Machinery (ACM).

Downloaded on 26 Apr 2023 17:03:57 SGT

Hardware-Software Collaborative Thermal Sensing in Optical Network-on-Chip-based Manycore Systems

MENGQUAN LI, Nanyang Technological University, Singapore and Chongqing University, China
WEICHEN LIU, Nanyang Technological University, Singapore
NAN GUAN, The Hong Kong Polytechnic University, Hong Kong
YIYUAN XIE, Southwest University, China
YAOYAO YE, Shanghai Jiao Tong University, China

Continuous technology scaling in manycore systems leads to severe overheating issues. To guarantee system reliability, it is critical to accurately yet efficiently monitor runtime temperature distribution for effective chip thermal management. As an emerging communication architecture for new-generation manycore systems, optical network-on-chip (ONoC) satisfies the communication bandwidth and latency requirements with low power dissipation. Moreover, observation shows that it can be leveraged for runtime thermal sensing. In this article, we propose a brand-new on-chip thermal sensing approach for ONoC-based manycore systems by utilizing the intrinsic thermal sensitivity of optical devices and the inter-processor communications in ONoCs. It requires no extra hardware but utilizes existing optical devices in ONoCs and combines them with lightweight software computation in a hardware-software collaborative manner. The effectiveness of the our approach is validated both at the device level and the system level through professional photonic simulations. Evaluation results based on synthetic communication traces and realistic benchmarks show that our approach achieves an average temperature inaccuracy of only 0.6648 K compared to ground-truth values and is scalable to be applied for large-size ONoCs.

CCS Concepts: • **Hardware** → **Emerging optical and photonic technologies**; **Thermal issues**; *Network on chip*; • **Computer systems organization** → *Embedded systems*;

Additional Key Words and Phrases: Hardware/software co-design, optical network-on-chip, chip temperature monitoring, micro-ring resonators, embedded systems

This work is partially supported by NSFC 61772094, China, and NAP M4082282 and SUG M4082087 from Nanyang Technological University, Singapore.

Authors' addresses: M. Li, Nanyang Technological University, Singapore and Chongqing University, China; email: mengquan.li@ntu.edu.sg; W. Liu, Nanyang Technological University, Singapore; email: liu@ntu.edu.sg; N. Guan, The Hong Kong Polytechnic University, Hong Kong; email: csguannan@comp.polyu.edu.hk; Y. Xie, Southwest University, Chongqing, China; email: yyxie@swu.edu.cn; Y. Ye, Shanghai Jiao Tong University, Shanghai, China; email: yeyaoyao@sjtu.edu.cn.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

1 INTRODUCTION

Due to the rapid growth of power density and the limited advancement of heat dissipation techniques, manycore systems suffer from overheating significantly. Silicon chips even cannot be fully utilized as a result of the “dark silicon” problem [22, 27], where a fraction of processor cores have to become powered-off or underclocked to maintain safe chip temperature. Effective thermal management solutions based on accurate yet efficient on-chip thermal estimation are critical for manycore systems [22, 23, 28]. State-of-the-art on-chip temperature sensing techniques are typically implemented in either hardware or software. Hardware-based thermal sensors possess high efficiency and measurement accuracy [5, 6, 9, 20, 26, 35]. However, additional chip area and hardware cost are required. In contrast, software-based approaches [21, 32, 47, 48] require no extra hardware but typically use thermal models or simulations for temperature prediction. They suffer from either expensive computation or low accuracy. Therefore, it is still challenging for manycore systems to implement accurate and efficient on-chip thermal sensing with trivial overhead.

As an emerging intercommunication architecture of new-generation manycore systems, optical network-on-chip (ONoC) [7, 23] offers unique advantages of high bandwidth, low latency, and low power dissipation. Taking a two-dimensional (2D) mesh-based ONoC as an example, as shown in Figure 1(a), vertically on top of the processing layer where processor cores are located, a photonic network (P-net) provides optical links between any pair of communicating processors for message transmission. Due to the inability of photonic to perform logical processing, an electronic network (E-net) is also provided for the control. The P-net and E-net together constitute an ONoC architecture. We also provide the logical view of a tile on the right side of Figure 1(a).

According to Figure 1(a), as the core functional component in ONoCs, an optical router is comprised of three parts: an *optical transmitter*, an *optical receiver*, and an *optical switching network*. Two specific switching network designs are provided for nonblocking 5×5 optical routers in Figure 1(b) and (c). There are five ports in a router. The injection and ejection ports connect the transmitter and the receiver, respectively. The other four bidirectional ports, including the west, south, east, and north ports, connect to its neighbor routers and are termed as passthrough ports. *Micro-ring resonators (MRs)* and *optical waveguides*, respectively, performing the switching and the transmission of optical signals, compose the *optical switching network*. The MRs, functioning as wavelength-selective optical switches, are essential to implement high-level routing policies in ONoCs.

MRs are highly sensitive to ambient temperature variations. As shown in Figure 2(a), the resonant wavelength of a MR red-shifts with increasing temperatures and blue-shifts with decreasing ones. The undesired mismatch between the signal wavelength and the resonant wavelength of the MR will result in additional optical power loss. Figure 2(b) shows the simulation results in Section 6.1. The optical power loss of a MR increases *monotonically* with the ambient temperature increasing over a large range from 300 K to 400 K. The intrinsic thermal sensitivity of MRs renders them an attractive choice for temperature sensing. The well-modeled temperature dependence of MRs that indicates the relationships between the resonant wavelength (and the optical power loss) of a MR and its ambient temperature make them practical. Compared to traditional electronic sensors, MR-based optical thermal sensors have favorable properties of compact size, immunity to electromagnetic interference, and robustness against mechanical shock [26]. However, most MR-based temperature sensors operate by monitoring the resonant wavelength shift with an expensive, narrow linewidth tunable laser or fine resolution optical spectrum analyzer [18]; consequently, heavy hardware costs and extra chip area are introduced, which makes them impractical for on-chip temperature measurement.

In this article, we present a new hardware-software collaborative thermal sensing approach for ONoC-based manycore systems. By utilizing the intrinsic thermal sensitivity of existing MRs and

(a) Left: A schematic diagram of a mesh-based ONoC architecture. Two primary layers consisting of an electronic control layer (E-net, bottom) and a photonic data layer (P-net, top) are shown; Right: A tile in logical view.

Fig. 1. ONoC architecture and example optical routers.

(a) The resonance wavelen. of MRs shifts with temp. changes. (b) A rise in ambient temp. increases the power losses of MRs.

Fig. 2. The characteristics of MRs change with temperature.

the inter-processor communications in ONoCs, it can implement accurate yet efficient runtime temperature monitoring while requiring no additional hardware support. We first quantitatively model the thermal sensitivity of MRs and develop a basic thermal sensing (BTS) module leveraging the idle injection and ejection ports of a single optical router. As the injection or ejection ports of routers are often occupied by inter-processor communications at runtime, we further propose a collaborative thermal sensing (CTS) approach for on-chip thermal estimation by combining the BTS module with a lightweight software solution. The CTS repurposes the optical routers for thermal sensing and continues using them for their intended purpose of communication without interruption. Based on device- and circuit-level photonic simulators that are widely used by the nanophotonics community for design and verification, simulation results show the average prediction error of the thermal sensitivity model is only 0.4985 K. Integrating with the precise model,

Table 1. Comparison with the State-of-the-art Hardware-based Thermal Sensors in Standard CMOS

Category	BJTs		MOSFETs		ETFs	our work
	Aita et al. [5]	Lakdawala et al. [20]	Chen et al. [9]	Anand et al. [6]	Sonmez et al. [35]	
Technology	0.7 μ m CMOS	32 nm CMOS	0.35 μ m CMOS	65 nm CMOS	40 nm CMOS	–
Inaccuracy (K)	0.25 (3σ)	<5	−0.25–0.35	± 2.3	± 1.4	0.6648
Chip area (mm^2)	4.5	0.02	0.6	$0.0004 + 0.0042$	0.00165	zero
Energy (pJ)	12.5×10^6	1.6×10^6	18.4×10^6	3.4×10^3	2.5×10^{12}	$\sim$ pJ
Hardware cost	medium		high		low	free

evaluation on synthetic communication traces and realistic benchmarks verify the effectiveness of the proposed CTS approach with an average inaccuracy of only 0.6648 K compared to ground-truth values and is scalable to be applied for large-size ONoCs.

The rest of this article is structured as follows: Section 2 investigates the previous on-chip thermal estimation techniques and briefly introduces ONoC architecture; fundamentals of the thermal sensitivity of MRs are presented in Section 3; Sections 4 and 5 propose the BTS module and the CTS technique, respectively; performance evaluations are presented in Section 6, and Section 7 concludes this article.

2 BACKGROUND AND RELATED WORK

2.1 On-chip Thermal Estimation

On-chip temperature estimation can be implemented in either hardware or software. As shown in Table 1, hardware-based electronic thermal sensors can be divided into three categories based on their basic operating principles: BJTs, MOSFETs, and Electro-Thermal Filters (ETFs) [29]. Generally, ETFs achieve higher accuracy and lower cost compared to transistor-based sensors (including BJTs and MOSFETs) thanks to the immunity to leakage current, process spread, and the mechanical stress of packaging [37]. While the transistor-based sensors have advantages in high efficiency and low power dissipation. Compared to BJTs, the temperature dependence of MOSFETs is well modeled and have lower energy consumption, but it involves higher costs for calibration than do BJTs. *In addition to electronic sensors*, optical thermal sensors using MRs have been studied recently. These sensors achieve high accuracy by coating the cladding with strong thermal-optical effect or employing dedicated MR architectures [18]. They are compact, immunity to electromagnetic interference, and robustness against mechanical shock and humidity. However, most of them operate by monitoring the resonant wavelength shift with an expensive, narrow linewidth tunable laser or fine resolution optical spectrum analyzer, which similarly introduces heavy overheads of chip area and hardware cost.

In contrast, Table 2 lists several state-of-the-art software-based approaches. ANSYS [4] is a commercial finite-difference and finite-element method (FEM) tool for thermal analysis. It achieves ultra-high accuracy but is computationally expensive due to the very fine-grained simulation grid. To achieve an acceptable estimation accuracy, ANSYS generally includes many more nodes in the thermal grid model than the number of processors. For example, the basic requirement for a 4×4 multicore system is a thermal grid with a 40×40 grid size. The node size of the thermal grid model in ANSYS simulator is 1,600, which is 100 times the number of processors. To reduce the computation time, analytical thermal models have been developed based on less grid nodes. There are three typical analytical models: *HotSpot*, *MatEx*, and the power blurring (PB) method, which sacrifice measurement accuracy for an acceptable time overhead. In addition, hybrid simulators

Table 2. Comparison with the State-of-the-art Software-based Temperature Estimation Techniques

Category	Numerical simulator	Analytical models			Hybrid Simulator	our work
	ANSYS® [4]	HotSpot 6.0 [47]	MatEx 1.0 [32]	Power Blurring [48]	NUMANA [21]	
Time Complexity	$O(N^2)$	$O(N^c)$ $c \in [1.5, 2]$	$O(N^2)$	$O(N \log(N))$	$\sim [O(N \log(N)), O(N^2)]$	$O(M^2)$ $M \ll N$
Max. Error	$\sim 0\%$	3.41%	$\sim 25.7\%$	13.7%	1.84%	–
Avg. Error	$\sim 0\%$	0.90%	$\sim 6.5\%$	2.5%	0.54%	–
Abs. Err. range (K)	$\sim 0-0$	$-1.4-1.15$	$\sim 0-4.2$	–	–	0.4499–1.7132
Avg. Abs. Err. (K)	~ 0	0.43	–	–	–	0.6648

have been proposed that achieve an estimation accuracy and a computation efficiency in between. As the relative and absolute errors of the analytical models and the hybrid simulator are compared with the error of ANSYS, we approximate the error of ANSYS as zero for simplicity.

These designs can be applied in both electrical and optical on-chip networks [22, 25, 29]. In this article, we consider the ONoC as target platform, because it is a promising solution for next-generation manycore systems. On top of ONoC-based manycore systems, our approach implements on-chip thermal sensing almost for free by utilizing the existing optical devices in ONoCs. It reduces the area and hardware overheads required for deploying other sensor designs and achieves high accuracy, high efficiency, and low power consumption with lightweight software computation, as shown in the last columns of Table 1 and 2. Moreover, in virtue of low overhead, this approach can also be used as an efficient complement of these existing thermal sensing techniques. Note that N and M in Table 2 are the node size of the thermal grid model in simulators and the size of the subset of routers in manycore systems, respectively. M is far less than N in general. Detailed performance analysis on our approach is shown in Section 6.

On-chip temperature information is crucial for system-level thermal management techniques. There are three kinds of application scenarios in ONoC-based manycore systems. First, it can be used for thermal-aware routing in ONoCs [23]. Due to thermal effects in ONoCs, chip temperature variations resulted from uneven power density and limited cooling techniques would cause significant optical power loss, which may counteract the power advantages of ONoCs. Thermal-aware routing techniques are critical for ONoCs to optimize communication energy efficiency in the presence of chip thermal variations. Second, runtime thermal management techniques such as workload migration and DVFS also require this information [22]. These techniques not only can reduce on-chip temperature gradients for mitigating the thermal effects in ONoCs but also can maintain chip thermal reliability by keeping every processor core within the safe temperature range. Last, to improve application performance, energy efficiency, and system reliability for thermal-safe guaranteed manycore systems, on-chip temperature information is also important for thermal-aware task mapping and scheduling [27, 28], through which system performance, energy efficiency, chip thermal reliability, and lifetime reliability can be systematically optimized.

2.2 ONoC Architecture

Due to the inability to perform inflight buffering and processing, an approach of *optical circuit switching* is applied in hybrid ONoC architectures: Optical communications are preceded by an electronic “*path-setup*” packet that is routed in the E-net to reserve an optical path; once the path is acquired, optical signals are transmitted end-to-end in the P-net. Generally, a communication link in the P-net is composed of an optical transmitter, an optical path and an optical receiver. The

Fig. 3. Micro-ring resonators.

transmitter converts electrical signals into optical signals (E-O conversion). A built-in microlaser source, such as VCSEL [36], can be implemented in the transmitter.¹ The VCSEL is connected with the underlying CMOS drivers using 3D integration technology and Through silicon Via (TSV) technique, similarly to the approach in Reference [15]. The output power of VCSELs is directly modulated by the driving current without optical modulators. Given the predefined driving current of VCSELs, we can easily know the initial optical power input into the link. The optical receiver uses high-resolution photodetectors (PDs) [34] to convert optical signals into electrical signals (O-E conversion), in which absorbed photons cause photo-induced carriers in the depletion region. A transimpedance amplifier (TIA) and a limiting amplifier (LA) are also included for current-to-voltage conversion and voltage amplification, respectively [19]. We can easily measure the received optical power using PDs. On the optical path between the transmitter and the receiver, multiple optical routers implement signal transmission. MRs, combined with optical waveguides, are used to perform switching operations in the routers. For every specific transmission, either no or only one MR is active in a router. Optical router designs follow this principle to minimize power loss for optimized communication performance and system reliability [16].

The intrinsic thermal sensitivity of MRs has aroused great attention recently. Ye et al. [45] systematically modeled and quantitatively analyzed the thermal effects of MRs. Padmaraju et al. [31] surveyed the thermal effects on MR-based devices. Bogaerts et al. [8] concluded the loss contributions of a ring resonator. Xiao et al. [40] presented an analytical model to quantify the resonator losses at room temperature without considering its physical structure and waveguide cleavage facets. These studies provide foundation for the temperature sensing technique developed in this article.

To achieve accurate yet efficient on-chip temperature sensing for ONoC-based manycore systems, three critical problems should be addressed. First, how to precisely model the intrinsic thermal sensitivity of optical devices, so as to provide a good theoretical foundation for the following thermal sensing techniques. Second, how to explore a basic thermal sensor design using the thermal sensitivity of existing devices in ONoCs. In this way, no extra chip area or hardware overhead is introduced. The third and most crucial step is to develop an efficient runtime thermal sensing approach, which can implement thermal estimation using existing devices without interrupting the devices' intended purpose of communication. Solutions to these problems are proposed in the following sections.

3 THE THERMAL SENSITIVITY OF MRS

As versatile devices in ONoCs, there are two different MR designs shown in Figure 3: the parallel switching element (PSE) design and the crossing switching element (CSE) design. Both consist of one ring and two straight waveguides. Ideally, when a MR is configured to be switched *off*, the optical signal from the *Input* port will be delivered to the *Through* port without directional change

¹Using grating couplers for fiber coupling to SOI chip and on-chip modulators, off-chip lasers can also be employed in optical transmitters.

(passive switching). Otherwise, the optical signal is resonated into the ring and delivered to the *Drop* port, which achieves a directional change (active switching). The PSEs turn the optical signal by 180° , while the CSEs turn it by 90° .

A MR resonates with light whose single-pass phase shift is a multiple of 2π , formulated as

$$m \cdot \lambda_{MR} = 2\pi R \cdot n_{eff}, \quad (1)$$

where λ_{MR} is the vacuum resonant wavelength of the MR, R is the average bending radius of the ring, n_{eff} is the effective index of the resonator, and m is a positive integer.

The resonant wavelength of a MR is sensitive to temperature fluctuation. Given the initial resonant wavelength, λ_0 , at the nominal operating temperature (typically room temperature), T_0 , the relation between the resonant wavelength of a MR and its ambient temperature can be expressed as follows [45]:

$$\lambda_{MR} = \lambda_0 + \rho_{MR} \cdot (T - T_0), \quad (2)$$

where ρ_{MR} is the temperature-dependent resonant wavelength shift coefficient of the MR, $\rho_{MR} = (\lambda_0 \cdot \delta n_{eff})/n_g \cdot n_g$ is the group index of waveguides and approximately equals 4.63 at 1,550 nm. The thermo-optic coefficient of the effective refractive index, $\delta n_{eff} = dn_{eff}/dT$, is smaller than that of silicon refractive index, $dn_{si}/dT = (1.86 \pm 0.08) \times 10^{-4} (K^{-1})$ [11].

The power losses of a MR primarily consist of coupling loss, propagation loss, and bending loss [8]. The optical power of signals whose wavelength are within the resonant wavelength range gets resonated into the cavity of a MR, while the signals whose wavelength are beyond the matched range is filtered and lost, which is the main cause of coupling loss. Propagation loss and bending loss are mainly caused by sidewall surface roughness and the bending radii of waveguides, respectively. In this article, we consider all these power losses. According to traveling wave theory [40], the power transmission of the through-port ($\phi_{through}$) and the drop-port (ϕ_{drop}) around the resonant wavelength (λ_{MR}) at constant ambient temperature can be calculated as follows:

$$\phi_{through} = \frac{(\lambda_{in} - \lambda_{MR})^2 + \left(\frac{FSR}{4\pi}\right)^2 (\kappa_d^2 + \kappa_p^2 - \kappa_e^2)^2}{(\lambda_{in} - \lambda_{MR})^2 + \left(\frac{FSR}{4\pi}\right)^2 (\kappa_d^2 + \kappa_p^2 + \kappa_e^2)^2}, \quad (3a)$$

$$\phi_{drop} = \frac{4 \times \left(\frac{FSR}{4\pi}\right)^2 (\kappa_d^2 \times \kappa_e^2)}{(\lambda_{in} - \lambda_{MR})^2 + \left(\frac{FSR}{4\pi}\right)^2 (\kappa_d^2 + \kappa_p^2 + \kappa_e^2)^2}, \quad (3b)$$

where Free Spectral Range (FSR) is the wavelength spacing between two resonance peaks of the MR. κ_e^2 and κ_d^2 are the fractions of optical power coupled into the ring from the input waveguide and the drop waveguide out of the ring, respectively. The relation $\kappa_e^2 = \kappa_d^2$ is generally recognized in symmetrically coupled add-drop MRs; consequently, we denote them uniformly as κ^2 hereinafter. κ_p^2 is the fraction of intrinsic power losses (such as bending, absorption, and surface scattering due to roughness) per round-trip in the ring. λ_{in} is the wavelength of input signals. According to Equation (3b), the -3-dB bandwidth of the drop-port power transfer spectrum can be expressed as $\theta = (FSR/2\pi)(2\kappa^2 + \kappa_p^2)$. Denoted γ_t as the minimum power transmission in the through-port, we can obtain $\kappa_p^2 = 2\pi \times \theta \sqrt{\gamma_t}/FSR$ when $\phi_{through} = \gamma_t$ at $\lambda_{in} = \lambda_{MR}$. The waveguide power coupling coefficients can be written as $\kappa^2 = \pi \times \theta(1 - \sqrt{\gamma_t})/FSR$.

Based on Equation (3), we formulate the thermal sensitivity of MRs as Equation (4), which describes the relationship between the optical power loss of an ON-state MR at the drop-port and its ambient temperature, T . Theoretical analysis and experimental results [41] show that the optical

Fig. 4. (a) The basic thermal sensing (BTS) module design. (b) The implementation of the BTS module in generic optical routers.

power losses at the drop-port of an *ON*-state PSE and an *ON*-state CSE that have identical physical parameter settings are approximately linearly dependent, with a fixed coefficient k ($k \neq 1$). We define the power loss of the CSE as ΔP and then that of the PSE is $k \times \Delta P$,

$$T = T_0 + \left(\frac{\sqrt{\alpha \cdot c^{(k \times \Delta P)} - \beta} + (\lambda_{in} - \lambda_0)}{\rho_{MR}} \right), \quad (4)$$

where $c = \sqrt[3]{10}$, $\beta = \theta^2/4$, and $\alpha = \beta \cdot (2\kappa^2/(2\kappa^2 + \kappa_p^2))^2$. In this article, we use the *ON*- or *OFF*-state to show the nominal states of MRs while use *active* or *passive* to describe their actual states. To perform on-chip communications, some of the MRs in optical routers will be activated, and they are designated to be in the *ON*-state while the others are passive and are designated in the *OFF*-state. It is noteworthy that the nominally *ON*-state MRs may actually be passive under large thermal variations, resulting in large optical power loss at the drop-port. By measuring the high power loss of the MRs resulted from thermal variations, we can estimate their ambient temperatures.

With a single-wavelength laser source whose output wavelength is equal to the nominal resonant wavelength of a MR (i.e., $\lambda_{in} = \lambda_0$), Equation (4) can be simplified to $T = T_0 + f(\Delta P)$ with all parameters constant for the MR. The ambient temperature of a MR can be obtained once we know the power loss of that MR. This physical model of MRs functions as the theoretical foundation of this article.

Equation (2) describes the relation between the resonant wavelength of a MR and its ambient temperature. Equation (4) well models the relation between the optical power loss of the MR and the temperature. Both of them can be applied for thermal sensing. To use Equation (2), it typically requires an expensive, narrow linewidth tunable laser or fine resolution optical spectrum analyzer to monitor the wavelength shift of the MR, which incurs heavy overheads. By contrast, measuring the power loss of the MR using PDs that are readily available in ONoCs, it requires no additional hardware or chip area when using (4), which is more suitable for on-chip thermal sensing in ONoC-based manycore systems.

4 BASIC THERMAL SENSING MODULE

The formula (4) well models the temperature dependence of MRs, which makes it possible to utilize single MR for on-chip thermal sensing. Motivated by it, we develop a basic thermal sensing (BTS) module for ONoCs in this section.

As shown in Figure 4(a), we can obtain the optical power loss of a MR as follows:

$$\Delta P_{MR} = P_{inj} - P_{ej}, \quad (5)$$

where ΔP_{MR} is the power loss of the MR; P_{inj} is the input power sent by the laser source and is typically known in ONoCs with the predefined driving current; P_{ej} is the received power measured by the PD at the receiver side. Based on it, the temperature of the MR can be derived from Equation (4).

Fig. 5. Chip thermal profiles of different task mapping patterns in the dark silicon era [27].

The BTS module can be employed as an independent thermal sensing module and can be placed in any area on chip requiring temperature monitoring at the expense of extra hardware and chip area. To achieve temperature sensing on ONoCs without requiring additional hardware support, we further customize the BTS design and implement it in optical routers. As shown in Figure 4(b), which is a local enlarged drawing of Figure 1(c), all the optical devices are readily available in typical ONoCs. By constructing an optical path from the idle injection port to the idle ejection port, passing through one MR, we can obtain the power loss of the MR based on the known sent power and measured received power; consequently, the temperature of the MR can be obtained using Equation (4).

Assumed that the heat in optical routers is evenly distributed due to the small footprint of routers, we can estimate the temperature of a router once the temperature of a MR in the router is known. In addition, we consider ONoCs based on 2D-mesh topology in this article, where optical routers are neatly placed at the top surface of the chip, evenly distributed. The on-chip thermal distribution can be estimated from the router temperatures with fine-grained consideration of the router-to-chip temperature offset. Moreover, optical routers are suitable to be used for on-chip temperature sensing, because they are evenly distributed across the network. It facilitates the runtime temperature monitoring of different chip regions based on the fine granularity of one sensor-per-router. These are important observations, which indicate that the on-chip temperature distribution can be estimated as long as the temperature of any MR in each optical router is obtained.

5 COLLABORATIVE THERMAL SENSING

Motivated: Using the BTS module, we can obtain the temperatures of a subset of optical routers whose injection and ejection ports are both idle. It is suitable for light-loaded ONoCs. However, for most of routers, their injection or ejection ports are often occupied by inter-processor communications at runtime. We can observe that (i) continuous technology scaling leads to a utilization wall challenge in MPSoCs, the rising “dark silicon” problem. Borrowed from Reference [27], the chip thermal profiles resulted from different dark silicon patterns are significantly different, as illustrated in Figure 5. Even under typical operating conditions, the steady-state chip temperature varies by 30 K across the chip [45]. In the dark silicon era, it would be difficult to predict the temperatures of processor cores based on their spatial correlation, because even neighboring cores may potentially have distinctly various temperature values. (ii) Optical routers are intended for inter-processor communications. It is not allowed to interrupt the normal communications when reusing the routers to perform thermal sensing. Unlike electronic channels where electronic packets can be buffered in situ, if an optical path is interrupted, then the source node has to send a

Fig. 6. The logical view of a communication path.

Fig. 7. An example of one communication path.

control packet to teardown the path, and a small acknowledgment packet may be sent from the destination node to the source for guaranteed delivery. The optical path is then reestablished and the bulk data have to be resent. Such an interruption operation introduces heavy overheads. (iii) It is impractical to wait for the communications to complete when measuring on-chip temperature distribution, especially under complex communication pattern and heavy network traffic. The execution of the BTS approach relies on the communication pattern and network traffic in ONoCs, which are various when executing different realistic applications. Considering the large volume of data transmitted on ONoCs, it typically takes a relatively long time to complete communications. However, the chip thermal profile in manycore systems varies according to the voltage/frequency level of processor cores and the heat dissipation efficiency between cores and heatsinks. It is ever-changing as a result of fine-granularity DVFS control. To sum up, it is not yet enough to achieve runtime thermal sensing using the BTS module alone.

To address this problem, we further propose a collaborative thermal sensing (CTS) approach. By combining the BTS module with a lightweight software solution, the CTS approach can repurpose the optical routers for runtime temperature sensing and continue using them for their intended purpose of communication without interruption. Not only in light-loaded ONoCs, it is also applicable to the applications with heavy-loaded communication demands, such as cloud computing. It supports fine-grained on-chip thermal sensing, which facilitates the deep optimization of manycore systems.

5.1 Loss Distribution of a Communication

Analyzing the routing path of a communication, based on the network topology and the optical router design, we can identify the input and output ports of each router used and then determine the active MR employed for each router in this path. Therefore, every communication path can be simplified as Figure 6.

Taking Figure 7 as a specific example, there will be two or more points with measurable power loss for a communication path. These points include the data injection point (the sender), the data ejection point (the receiver), the through points (if any), the turning points (the routers where the signal direction changes), and the waveguides (for signal propagation). Based on the known input power sent by the sender (denoted by P_{inj}) and the power of optical signals measured by the PD at the receiver side (denoted by P_{ej}), according to the energy conservation law, the power loss for a path can be formulated as follows:

$$\Delta P_{inj} + \sum \Delta P_{through} + \sum \Delta P_{turn} + \sum \Delta P_{wg} + \Delta P_{ej} = P_{inj} - P_{ej}, \quad (6)$$

Fig. 8. An example of the CTS technique.

where the variables ΔP_{inj} , $\Delta P_{through}$, ΔP_{turn} , ΔP_{wg} , and ΔP_{ej} are the optical power losses for the sender, the through point(s), the turning point(s), the waveguide(s), and the receiver, respectively.

5.2 Collaborative Thermal Sensing Matrix

Every router may participate in multiple data communications and play multiple roles simultaneously including sender, receiver, and(or) intermediate node. There are typically multiple communications in ONoCs. By constructing the linear equations (6) of the multiple communications, we obtain a linear equation system.

We illustrate this with an example of a 3×3 2D-mesh ONoC equipped with Crux routers, shown in Figure 8(a). We obtain the linear equations of the three data communication paths (DPs) as:

$$\begin{cases} \Delta P_2 + \Delta P_1 + \Delta P_4 = P_{inj} - P_{ej}^4 - 2 \times \Delta P_{wg} \\ k \times \Delta P_5 + \Delta P_6 = P_{inj} - P_{ej}^6 - \Delta P_{wg} \\ \Delta P_9 + \Delta P_8 + k \times \Delta P_5 = P_{inj} - P_{ej}^5 - 2 \times \Delta P_{wg} \end{cases}, \quad (7)$$

where P_{ej}^i denotes the optical power received by the destination router i for each path. P_{inj} is the input power sent by the transmitter in every path. ΔP_{wg} is the power loss of a straight waveguide connecting adjacent routers. It is determined by sidewall surface roughness and waveguide length and can be modeled as $\Delta P_{wg} = \varepsilon \cdot Len$, where ε is a constant mainly determined by the residual surface roughness on the etched sidewalls of the waveguide and Len is the waveguide length. Simulation results show that the temperature variation has a trivial impact on waveguide loss [38]. Therefore, under typical silicon photonic fabrication technology, ΔP_{wg} is constant for the waveguides of the same length. The first and the third paths contain two waveguides and thus a factor of 2 should be multiplied. By using post-fabrication calibration [43] that is widely employed for typical ONoCs to overcome technology variances, all the CSEs in an ONoC are considered to be identical and so do all the PSEs. Assuming the CSEs and the PSEs have the same physical parameter settings, the optical power loss of the CSEs and the PSEs in router i are ΔP_i and $k \times \Delta P_i$, respectively. Considering the constant coefficient k , there is only one unknown variable, ΔP_i , in the router i . In this example, R_5 sends data to R_6 with an injection-to-east transmission and receives data from R_8 with a south-to-ejection transmission, both of which involve PSEs as shown in Figure 1(c); consequently, the factor k should be applied.

The set of Equations (7) can be generalized as matrix and vector multiplications:

$$\overline{A} \times \overline{\Delta P} = \overline{L^{path}}$$

where $\overline{\Delta P}$ is the vector of the power loss for all the routers, $\overline{L^{path}}$ is the vector of the total power loss of routers for all the paths, and $\overline{A}$ is the coefficient matrix,

$$\overline{A} = \begin{pmatrix} 1 & 1 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & k & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & k & 0 & 0 & 1 & 1 \\ & & \dots\dots & & \mathbf{0} & & \dots\dots & & \end{pmatrix}$$

$$\overline{\Delta P} = (\Delta P_1, \Delta P_2, \Delta P_3, \Delta P_4, \Delta P_5, \Delta P_6, \Delta P_7, \Delta P_8, \Delta P_9)^T$$

$$\overline{L^{path}} = (P_{inj} - P_{ej}^4 - 2\Delta P_{wg}, P_{inj} - P_{ej}^6 - \Delta P_{wg}, P_{inj} - P_{ej}^5 - 2\Delta P_{wg}, 0, 0, 0, 0, 0, 0)^T.$$

The first three rows of $\overline{A}$ model the three paths, while the rest of the matrix is filled with zeros.

To obtain a unique solution for the linear equations with N variables, the coefficient square matrix $\overline{A}$ must be nonsingular. In our example, there are seven variables but only three equations; consequently, $\overline{A}$ cannot be full rank. To make the problem solvable, we can establish more equations by constructing some auxiliary paths (APs). Note that we merely need to model the working routers whose injection or ejection ports are occupied by data communications, because the temperatures of the free routers whose injection and ejection ports are both idle can be obtained using the BTS module in Section 4. This is an important observation that can significantly reduce the problem size (from the size of the entire network to the size of the set of working routers). In this example, the free routers R_3 and R_7 in Figure 8(a) can be excluded from consideration. Thus, the size of the coefficient matrix $\overline{A}$ is reduced to 7×7 .

THEOREM 5.1. *Given an ONoC subset with N routers and M DPs, where all the M paths use only the N routers including data injection, data ejection, and data passthrough, we can always construct $N - M$ APs by using the N routers.*

PROOF. We classify the ports of an optical router into *incoming ports* (including the injection and passthrough ports) and *outgoing ports* (including the ejection and passthrough ports), through which the optical signals transmitted into and out of the router, respectively. Incoming ports and outgoing ports exist in pairs in a router and must be used in pairs as well due to the inability of routers to buffer optical signals. Start with a router that has an idle injection port p_i , $p_i \in \text{incoming ports}$, there must exist an idle port p_j , $p_j \in \text{outgoing ports}$, in the router. We can follow the outgoing port p_j to either the local processor (p_j is the ejection port) or to the next router (p_j is a passthrough port). In the latter case, we continue searching for an outgoing port in the next router and repeat the process until an available ejection port is found, at which point an AP is constructed. Similar rules can always be applied to construct the rest of the APs. Because N routers have N injection and ejection ports and M DPs use exactly M ports of them, we can always construct the rest $N - M$ APs. $\square$

Theorem 5.1 guarantees that N equations can be formulated for N routers. Furthermore, all the row vectors of $\overline{A}$ must be linearly independent to solve the linear equations. When a new AP is constructed, we need to analyze whether it is linearly independent of the existing paths to confirm its usefulness in resolving the set of linear equations. We call two paths linearly independent iff the coefficient vectors of the two paths are linearly independent. *If adding an auxiliary path increases the rank of $\overline{A}$ by 1, then the auxiliary path is linearly independent with the existing paths and should be accepted.* This evaluation method can be applied in the auxiliary path construction algorithm to quickly determine the linear independence among communication paths.

5.3 Constructing Auxiliary Paths

Given a set of DPs, the number of AP alternatives is exponential. Indeed, the problem of constructing a set of APs to make the equations solvable is NP-hard. Based on the analysis above, we present a heuristic algorithm to construct feasible auxiliary paths efficiently.

As shown in Algorithm 1, given an ONoC and a set of DPs, we can identify a subgraph of the ONoC containing only the working routers and the available links among them (Lines 1–3). These resources are utilized to construct APs. First, the BTS approach proposed in Section 4 can be applied for the routers with both injection and ejection ports available but some passthrough ports occupied (Lines 4–6). Second, we construct an AP from each sending router whose injection port is available. Along the available links among routers, we can find a path from the sender to a receiving router with the ejection port available, by Depth First Search (DFS) (Lines 7–18). If the rank of the coefficient matrix $\bar{A}$ is increased by 1, then the AP is accepted (Lines 14–16). Otherwise, the AP is rejected, and DFS continues to the next router until a feasible AP has been constructed. If no such path can be constructed, then the algorithm returns a failure indication (Lines 17 and 18). Note that the function `DepthFirstSearchNextRouter()` can move back directly to the sender after the branch along one direction has been visited (four directions at most) (Line 10). We illustrate this algorithm in Figure 8(b). The free routers R_3 and R_7 are excluded from our equation system as mentioned above. First, we apply the BTS approach for the working routers R_1 and R_8 whose injection and ejection ports are both idle but some passthrough ports are occupied. Then, we construct two APs, (R_4, R_5, R_8, R_9) and (R_6, R_5, R_2) , starting from R_4 and R_6 , respectively.

ALGORITHM 1: Auxiliary path construction algorithm.

Data: The *ONoC*, the set of *DPs*

Output: Success/Failure, the constructed set of *APs*

```

1   $N = \text{GetTheNumOfWorkingRouter}(\text{ONoC}, \text{DPs});$ 
2   $A = \text{Path2Matrix}(\text{DPs});$ 
3   $\text{APs} = \emptyset;$ 
4  while  $\text{router} = \text{GetRouterWithBothIdleInjEjPorts}()$  do
5       $\text{ap} = \text{router}; \text{APs} += \text{ap};$ 
6       $A = A + \text{Path2Matrix}(\text{ap});$ 
7  while  $\text{Rank}(A) < N$  do
8       $\text{sender} = \text{GetRouterWithOnlyIdleInjPort}();$ 
9       $\text{ap} = \text{sender}; \text{found} = \text{false};$ 
10     while  $\text{router} = \text{DepthFirstSearchNextRouter}(\text{ap})$  do
11          $\text{ap} += \text{router};$ 
12         if  $\text{router}$  is a receiver then
13              $A' = A + \text{Path2Matrix}(\text{ap});$ 
14             if  $\text{Rank}(A') == \text{Rank}(A) + 1$  then
15                  $A = A'; \text{APs} += \text{ap}; \text{found} = \text{true};$ 
16                 Break;
17     if  $\text{!found}$  then
18         return Failure;
19 return (Success,  $\text{APs}$ );
```

Fig. 9. The theoretical and FDTD-simulated responses of the MR at room temperature.

Complexity analysis: The complexity of Algorithm 1 is $O(N^2)$, where N is the number of the working routers. Each router has a fixed number of five ports. The complexity of constructing an AP is $O(N)$ using the depth-first search without backtracking. With N APs to be constructed in the worst case, this algorithm has a polynomial time complexity of $O(N^2)$. Furthermore, by reducing the search space from the entire network to the set of working routers, the execution time of Algorithm 1 is largely reduced in practice use.

Although we consider 2D-mesh based ONoCs as the target platforms in this article, Algorithm 1 can be extended to be applied in other network topologies, such as torus. The difference is, as each router in torus-based networks has four neighbors, there are four candidates when finding the next router for a router using DFS in Line 10; while in mesh-based networks, the routers located at the edges of the networks have two or three neighbors, thus there are only two or three candidates when finding the next router for these routers. In addition, we do not restrict the routing algorithm for APs to obtain high flexibility; that is, more AP alternatives can be attempted to be constructed. The APs would adapt to the DPs without interrupting their execution, and only a small amount of optical data is transmitted along every AP for energy saving. The data communications have priority over auxiliary communications. If the resources that are being occupied by APs for thermal sensing are required by new data communications, then the APs would be released immediately with tiny overhead, and the available resources can be used for the new data communications. Based on the given DPs, the APs are constructed in batch without deadlock by using the centralized algorithm (Algorithm 1). Besides, considering the hybrid network architecture employed in our work, the path establishment is performed via the electrical network. Deadlock in the electrical network can be avoided using virtual channel flow control [25], and the photonic network is inherently deadlock-free due to circuit switching and predetermined routing path.

Based on the constructed APs using Algorithm 1, the linear equations can be solved with a unique solution. Combined with the BTS module proposed in Section 4, the optical power losses for all the routers are known. We can estimate their temperatures (as well as the chip thermal distribution) using Equation (4).

6 EVALUATION

6.1 Accuracy of the Thermal Sensitivity Model

We use a compact silicon-on-insulator (SOI) MR with a radius of $4\text{ }\mu\text{m}$ to experimentally verify the accuracy of our thermal sensitivity model (Equation (4)). The coupling gaps between the silicon waveguides and the ring are 100 nm with a cross section of $400\text{ nm}\times 180\text{ nm}$. We simulate the MR and study its optical characteristics through device-level 3D FDTD simulations [1]. Figure 9(a) shows a perspective view of the MR in FDTD simulator. FDTD method is a recognized numerical

Table 3. Loss Parameters for the MR

λ_{res} (nm)	1505.28	1525.61	1546.29	1567.54	1589.59
FSR (nm)	20.33	20.51	20.97	21.65	22.05
γ_t	0.01146	0.01302	0.00350	0.00883	0.00454
θ (nm)	1.03	1.16	1.40	1.92	2.53
κ_p^2	0.0339	0.0406	0.0248	0.0524	0.0486
$\kappa^2 = \kappa_e^2 = \kappa_d^2$	0.14158	0.15745	0.19737	0.25243	0.33618
Exp. loss (dB)	1.1564	0.8817	0.6694	0.8944	0.8799
Theo. loss (dB)	0.9834	1.0522	0.5300	0.8570	0.6057

analysis technique used for modeling computational electrodynamics by solving Maxwell's equations.

Supporting the fundamental TE mode for wavelengths at 1,500 nm to 1,600 nm, Table 3 lists all the resonance peaks λ_{res} and their optical responses (FSR, the minimum power transmission of the through-port γ_t and the -3-dB bandwidth of the drop-port θ) obtained by FDTD simulations at room temperature (300 K). We compare the theoretical losses obtained by Equation (3) with the FDTD-simulated results. Small discrepancies (<0.28 dB) between them are observed due to the uncertain contributions from coupling and waveguide cleavage facets.

Using the extracted parameters in Table 3, we plot the theoretical response curves and the FDTD-simulated results at center resonant wavelengths of ~ 1526 nm and ~ 1546 nm at room temperature. In Figure 9(b) and (c), the blue and black dots are the responses of the MR at the through-port and drop-port in FDTD simulations, respectively, while the red ($\phi_{through}$) and green (ϕ_{drop}) curves are derived from Equation (3). It shows that the theoretical curves match well with simulation results. These results verify the high accuracy of our thermal sensitivity model at room temperature.

Given the incident light as the fundamental TE mode for the wavelength at ~ 1546 nm, with which the MR resonates at room temperature, we test the accuracy of the thermal sensitivity model under different ambient temperature from 300 K to 380 K through FDTD simulations. A temperature of 380 K is an extremely high value for silicon chips. But the our model is not only applicable to the MRs in ONOCs but also suitable to be used for dedicated MR-based thermal sensors. In every group of experiments, setting a simulation temperature for FDTD simulator, we can obtain the responses of the MR (e.g., the optical power loss in the drop-port) under this temperature from simulation results. For every simulation, on the basis of the temperature profile data obtained from a 3D heat transport simulator (called HEAT) [2] that is a finite element method providing designers with comprehensive thermal modeling capabilities, a recognized linear model is applied in the FDTD simulator to calculate the material refractive index under the simulation temperature, formulated as $n + iy = (n_0 + \Delta n) + i(y_0 + \Delta y)$, where n and y , respectively, are the real and imaginary parts of the complex refractive index of the material at simulation temperature T ; n_0 and y_0 are the real and imaginary parts of the unperturbed refractive index at default temperature T_0 (room temperature in this article), and $\Delta n = \frac{dn}{dT}(T - T_0)$ and $\Delta y = \frac{dy}{dT}(T - T_0)$ are the changes in the real and imaginary part of the refractive index, respectively. Using the linear model and the obtained temperature profile data from HEAT, the FDTD simulator can calculate the effective refractive index of the simulated MR and then obtain its optical characteristics under the simulation temperature.

As shown in Figure 10, the red curve is the response of the MR under different simulation temperature in FDTD simulator; the green curve is derived from Equation (4). An average difference of only 0.4985 K (a minimum absolute difference of 0.0021 K and a maximum difference of 1.7166 K)

Fig. 10. The accuracy of the thermal sensitivity model based on MRs of different structures.

exists between the simulation temperature (ground-truth value) and the temperature obtained from our thermal sensitivity model. This accurate thermal model provides a good foundation to support our CTS technique.

Our methodology and simulations can be conducted based on MRs of different structures. The MR simulated in this article is based on typical dimensions, similarly to the MR in Reference [42]. We use this MR to experimentally verify the effectiveness of our thermal model and methodology. To validate the applicability of the proposed model (Equation (4)) and our methodology on MRs of different structures, we have conducted another two groups of simulations. Generally, the extinction ratio (ER) of a MR is predetermined by the structure design. According to the simulation results above, the ER of the simulated MR is about 18 dB. Similar MRs with high ERs are employed in Reference [39, 46]. In the simulations, we use two MRs with different ERs. Their ERs are approximately 13 dB and 22 dB. As shown in Figure 10, the blue and yellow curves are the simulated results of the two MRs and the black and purple curves are derived from Equation (4). Results validate the prediction accuracy of our model based on the MRs of different structures, with the average errors of only 0.3201 K and 0.7321 K. Note that the constant parameters in our model will take different values for MRs of different structures (thus different ERs). For a specific MR, the values of the constant parameters in our model are fixed and will be determined by its optical responses (e.g., FSR, γ_t and θ shown in Table 3).

6.2 The Feasibility of the CTS Technique

Besides the accuracy of the thermal sensitivity model, the effectiveness of the CTS technique relies on whether it can always construct feasible APs to successfully solve the linear equations within polynomial time. To test the feasibility of the CTS technique, we consider 2D mesh-based ONoCs with size range from 2×2 to 18×18 as the target platforms. For every ONoC, we set nine discrete communication load rates from 0.1 to 0.9, with an interval of 0.1. The communication load rate is the ratio between the number of DPs and the maximum number of communications afforded by ONoCs; for example, the maximum communication load of a 2×2 ONoC is 4, since at most four communications exist simultaneously. We conduct 100 groups of experiments under each communication load rate. In every group of experiment, the sources and the destinations of the data communications among processors are randomly generated.

We count the percentage of the cases where the CTS technique finds a set of APs to successfully solve the linear equations and denote it as success rate. As shown in Figure 11(a), the CTS technique achieves an average success rate as high as 98.94% and is scalable to large-size ONoCs. A 100%

Fig. 11. The feasibility of the CTS technique.

success rate is achievable with additional efforts. It is observed that the success rates of the CTS in small-size ONoCs are lower than those in large-size ONoCs. That is because, for the small-size ONoCs with lightweight communication loads, the number of DPs and the number of AP alternatives are too less to find a set of feasible APs. For example, there is only one DP in a 2×2 ONoC when the communication load rate is below 50%; consequently, only one AP alternative exists and it is very possibly linearly dependent on the DP. Thus, this is not a fundamental limit of our CTS approach but a result of the experimental setup.

To further analyze the success rate of the CTS technique under different communication load rates, we extract the results on 3×3 , 9×9 , 14×14 , and 18×18 ONoCs. As shown in Figure 11(b), the CTS technique adapts well to heavily-loaded communication situations though its success rate varies slightly with different load rates.

6.3 The Effectiveness of the CTS Technique

Considering the impact of crosstalk noise and the insertion loss of optical devices that commonly exist in ONoCs, we further evaluate the effectiveness of the CTS technique at the system level based on both synthetic communication traces and realistic benchmarks.

The target platforms are 2D mesh-based ONoCs with sizes from 2×2 to 18×18 . We simulate the communications in ONoCs through a photonic integrated circuits (PICs) simulator, called INTERCONNECT [3]. Incorporating the resultant compact model parameters (e.g., S parameters) extracted from device-level simulations using FDTD simulators, INTERCONNECT can simulate PICs accurately in both the frequency and time domain and obtain the total circuit response using scattering data analysis. Figure 12(a) provides an example view of one communication path simulated in INTERCONNECT, including an optical transceiver (ONA), three optical routers (R_{ij}), and waveguides between the routers (W_{ij}). We also illustrate the internal structure of a router, in which the S parameters of switches are extracted from 3D FDTD simulations. Optical routers are connected by 1-mm-length SOI waveguides in the simulations. Experimental results show that the optical loss of the waveguides used in the simulations is approximately 0.2803 dB/mm.

To base simulations on synthetic communication traces, we conduct 10 groups of experiments for each size of ONoC. In every group of experiments, the source, destination and volume of data communications among processor cores are randomly generated. Based on the generated DPs, we construct a set of APs following Algorithm 1 and then simulate all the communications (including the DPs and the APs) on the ONoC through INTERCONNECT. The total power loss of each

(a) An example of one communication path in circuit-level simulations.

Fig. 12. The accuracy of the CTS technique based on (b) synthetic communication traces and (c) realistic benchmarks (error bars donate S.D).

communication path can be obtained from simulation results. The temperature of every individual router is also randomly generated with a range from 300 K to 380 K, which is set as the simulation temperature of the router in the simulator. By comparing the temperatures obtained by the CTS with the simulation temperatures (ground-truth values), we can obtain the measurement inaccuracy of the CTS technique. As shown in Figure 12(b), the CTS achieves high measurement accuracy with an average error of only 1.1003 K and is scalable to large-size ONoCs. As the ONoC size increases, the insertion loss of the growing number of optical devices and the crosstalk noise among them would aggravate the sensing accuracy of the CTS. We would fail to obtain accurate temperature values based on inaccurate loss measurements. Nevertheless, with the improvement of silicon photonic technologies, this kind of negative effect would be largely alleviated.

We further conduct evaluations based on a set of realistic applications, including *autocor*, *audiobeam*, *tde_pp*, *fmradio*, *filter bank*, and *beamformer* in StreamIt benchmarks [28]; *8_Rlattice* and *IIR filter* in the DSP-stone benchmark; and the industry standard H.264 HDTV decoder [44], to evaluate the CTS technique in practice use. We build a realistic simulator in Python to produce the data communication traces and thermal profiles for the execution of every application, both of which are the inputs to the CTS. First, the data communication traces are generated using the task mapping algorithm proposed in Reference [44]. Then, the power modeling of the task executions is performed by integrating McPAT v1.0 [24] into the Python simulator. Considering the DVFS capability of modern processors, McPAT models computation power under different voltage/frequency levels based on out-of-order Alpha 21346 cores in 22-nm technology. The generated power consumption traces of the processor cores are used to obtain steady-state chip thermal profiles through HotSpot v5.02 [17], which exhibits on-chip thermal distribution. Finally, by approximating router temperature to their adjacent processor core, using the generated data communication traces and the corresponding temperature profiles, we set the simulation temperatures for the routers and simulate the communications in ONoCs through INTERCONNECT, similar to the experiments based on the synthetic communication traces. In the experiments, we have

validated the accuracy of the CTS against HotSpot simulator, which is widely used in manycore systems for temperature prediction. For simplicity, we assume that the steady-state chip thermal profiles obtained by HotSpot are the ground-truth values in the experiments without regard to its errors, and obtain the monitoring error of our approach by comparing with them. Note that our approach is also applicable in actual use. By setting the real chip temperature distribution as the ground-truth values, our approach works in the same way as it works in the experiments. As shown in Figure 12(c), the average inaccuracy of the CTS technique is 0.6648 K. The evaluation results are consistent with those based on the synthetic communication traces, which validates the effectiveness of our CTS approach for thermal monitoring.

6.3.1 Overhead Analysis. Our CTS technique can be integrated into the centralized resource manager in the operating system (OS). The manager takes control of the whole network and knows the occupation of optical links. Based on the task-to-core mapping determined by a given task mapping algorithm, the inter-processor communications (i.e., DPs) are known by the manager. The known information (including the information of optical link occupation and DPs) is helpful for constructing auxiliary paths. When the OS receives a temperature-sensing request, it would construct a feasible set of APs based on the DPs by using Algorithm 1 if APs are needed. A set of AP-constructed tasks are delivered to the corresponding processor cores; consequently, the APs are established, along which small amounts of optical data are transmitted. After the receiver of a communication path (including DPs and APs) obtains the received optical power, a small packet containing that information will be sent to the manager via the electronic network. The total loss of every communication path can be easily obtained with the known sent power and measured received power. After solving the constructed linear equations, the OS can calculate the temperature of each router using Equation (4).

The overheads of the CTS technique are mainly incurred by (i) executing Algorithm 1, (ii) constructing APs, and (iii) solving the linear equations. The cost incurred by sending the AP-constructed tasks to the cores and the power information to the OS is a general issue in centralized thermal management. Similar overhead also exists when using traditional thermal sensors, where the centralized manager send the temperature-sensing command to the sensors and the temperature of each core is sent to the manager.

Time overhead: It requires no extra time for AP construction, because the processes of auxiliary communication and data communication are in parallel. Specifically, the latency of AP construction includes the latencies of the VCSELs and PDs, the path-setup latency in the electronic network and the data transmission latency in the photonic network. The VCSELs and PDs both have nanosecond or lower latency, which is at the same order with the path-setup latency in ONoC [7]. As the transmission rates of O-E interfaces reach tens of Gbps and the transmission velocity of light in silicon waveguides is $\sim 6.6 \times 10^7$ m/s, the time costed to transmit small amounts of data along optical paths is also typically at nanosecond level once the paths are established [33]. In summary, the time cost for AP construction is on the order of nanosecond. Algorithm 1 has a polynomial time complexity of $O(N^2)$, where N is the number of the working routers. By reducing the size of search space from the whole network to the set of working routers, its execution time is further reduced. Similarly, solving the set of linear equations can be implemented with high efficiency by using the Gaussian elimination method and LUP decomposition, which is also of polynomial time complexity [12].

Energy overhead: Besides the energy consumed by executing heuristic Algorithm 1 and the light-weight solver of linear equations, the energy required for AP construction is another contributor to the total energy consumption of the CTS. The VCSELs and PDs both consume 1 pJ/bit for O-E and E-O conversion [14]. From Reference [14], the energy required by path establishment in the

electronic network is formulated as:

$$E_{path-setup}^e = E_{int}^e \cdot L_{ctrl}^e \cdot h + E_{cu}^e \cdot (h + 1), \quad (8)$$

where E_{int}^e is the average energy required to transfer a single bit through electrical interconnections, L_{ctrl}^e is the total size of the control packets, h denotes the number of hops from the source to the destination in the electronic network, and E_{cu}^e denotes the average energy required by the control unit to make a decision for a single packet (generally 1.5 pJ). A typical control packet contains 8+1 bits and requires 0.52 pJ/bit energy in an electronic network [14]. The energy consumed for data transmission in the photonic network is extremely low thanks to the property of *bit-rate transparency*—the energy dissipation of photonic switches does not scale with the bit rate, because they switch on and off per packet instead of per bit of the transmitted data. To sum up, the total energy consumption of our approach is on the order of pJ. Moreover, as the power consumption of the CTS is kept at a negligible level, which is much lower than the requirement (i.e., at nJ level [10]), the errors due to self-heating are largely mitigated.

6.3.2 Effect of Process Variations. Fabrication-induced process variations (PVs) in practice is a device-level limitation in ONoCs. For a MR, a variation of 1 nm in waveguide width introduces approximately 0.58-nm resonant wavelength shift [30]. Recent fabrication results show that the physical dimension variation among MRs is observed as 0.37 nm within a die [43], resulting in a resonant wavelength drift of 0.2146 nm. According to Equation (2), given the temperature-dependent resonant wavelength shift coefficient, ρ_{MR} , we can easily obtain the relationship between the resonant wavelength shift of a MR and its temperature variation. We have experimentally tested the value of ρ_{MR} based on the simulated MR through 3D FDTD simulations. Simulation results show that ρ_{MR} is approximately 0.0658 nm/K at the 1,550 nm wavelength range. Similar results are presented in References [23, 26]. Consequently, within a die, the PVs among MRs would cause a temperature error of about ± 3.2614 K in our technique, which is acceptable for the current chip thermal management techniques [17]. Besides, the accuracy degradation for temperature monitoring due to PVs would be reduced continuously as the nanophotonic technology advances.

Most of PVs are static and can be calibrated during the chip testing stage. Post-fabrication calibration techniques are widely employed for MRs to precisely correct the wavelength drifts due to PVs [30, 43]. The impact of PVs on our technique can be mitigated. Furthermore, we focus on implementing chip-wide temperature sensing by utilizing the intrinsic thermal sensitivity of MRs. A system-level thermal measurement methodology is proposed. State-of-the-art techniques addressing PVs at the device level are orthogonal to our approach and are considered good complement to complete it.

However, based on the PV maps obtained during chip testing, we can also conduct detailed statistical analysis (such as Monte Carlo analysis) to simulate a different optical power value, $\Delta P'$, for every MR instance laid out on the chip. Assumed that the optical power loss of a standard MR (without PVs) is ΔP under ambient temperature T , for any MR instance, its loss value $\Delta P'$ approximates $k' \times \Delta P$, formally $\Delta P' = k' \times \Delta P$, k' is different for MR instances. In this way, Equation (7) can be formulated as follows:

$$\begin{cases} k'_1 \times \Delta P_2 + k'_2 \times \Delta P_1 + k'_3 \times \Delta P_4 = P_{inj} - P_{ej}^4 - 2 \times \Delta P_{wg} \\ k'_4 \times \Delta P_5 + k'_5 \times \Delta P_6 = P_{inj} - P_{ej}^6 - \Delta P_{wg} \\ k'_6 \times \Delta P_9 + k'_7 \times \Delta P_8 + k'_8 \times \Delta P_5 = P_{inj} - P_{ej}^5 - 2 \times \Delta P_{wg} \end{cases}, \quad (9)$$

where ΔP_j is the optical power loss of the standard MR in router j . k'_i ($i = 1, 2, \dots, 8$) is the PV-induced coefficient that is obtained by dividing the loss of the active MR (under PVs) in router j by that of the standard MR in this router; thus, $k'_i \times \Delta P_j$ is the optical power loss of the active MR

Fig. 13. (a) The effect of device-level wavelength tuning on MRs. (b) Our technique is compatible with local wavelength tuning technique.

in router j . Note that we do not need to distinguish PSEs and CSEs in this case. According to the PV maps obtained during chip testing, the values of k'_i ($i = 1, 2, \dots, 8$) are known. Therefore, if solving the linear equation system successfully, then we can know the optical power loss of the standard MR in every router; consequently, the temperature value of the router can be estimated using Equation (4). Our thermal sensing methodology is still feasible. In addition, considering the PVs among MRs, the coefficient square matrix $\bar{A}'$ is expressed as

$$\bar{A}' = \begin{pmatrix} k'_2 & k'_1 & 0 & k'_3 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & k'_4 & k'_5 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & k'_8 & 0 & 0 & k'_7 & k'_6 \\ \dots & \dots & \dots & 0 & \dots & \dots & \dots & \dots & \dots \end{pmatrix}.$$

Compared to the matrix $\bar{A}$ where only two coefficients (1 and k) exist, $\bar{A}'$ is easier to be full rank, which potentially increases the success rate of the CTS technique and thus enhances its feasibility.

6.3.3 Discussion on Device-level Wavelength Tuning Technique. Device-level wavelength tuning is one of the common techniques that can be applied for MRs to implement reliable inter-processor communication in ONoCs [23, 31]. In this section, we analyze the applicability of our technique under the circumstances where the local wavelength tuning technique is applied. Figure 13(a) illustrates the effect of local wavelength tuning on MRs. Due to thermal variations, the resonant wavelength of a MR red-shifts, termed as TV-drift (e.g., from the black wave to the red wave in Figure 13(a)). Using local microheaters and auxiliary PDs, wavelength tuning technique compensates for the wavelength shift (e.g., from the red wave to the grey wave in Figure 13(a)) and dynamically maintains the resonance of the MR throughout the duration of its operation at the expense of extra regulation power consumption. As the resonant wavelength is almost realigned with the nominal one, the optical power loss of the MR caused by the undesired TV-drift would be largely reduced.

Although we do not consider local wavelength tuning in this article, it is compatible with our thermal sensing methodology. Given an ONoC where local wavelength tuning is employed for MRs, the losses of the MRs obtained by the CTS technique are the losses after tuning but not those before wavelength tuning; consequently, the thermal sensitivity model (i.e., Equation (4)) is no longer applicable. Nevertheless, we can easily obtain the loss of a MR before tuning based on the regulation power consumed by wavelength tuning and the loss of the MR after tuning obtained by the CTS.

As shown in Figure 13(b), the top-down process shows the flow of local wavelength tuning technique. If operating from the bottom up, then the optical power loss of the MR before tuning (ΔP) can be obtained, which is the sum of the loss after tuning (ΔP_{tuned}) and the loss reduced by the wavelength tuning technique ($\Delta P_{reduced_by_WT}$). Formally, $\Delta P = \Delta P_{tuned} + \Delta P_{reduced_by_WT}$. Analyzed in Reference [26], the power loss reduced by wavelength tuning can be expressed as

$\Delta P_{reduced_by_WT} = 10 \log \left(\left(\frac{2\kappa^2 + \kappa_p^2}{2\kappa^2} \right)^2 \left(1 + \frac{4P_{WT}^2}{\epsilon^2 \theta^2} \right) \right)$, where P_{WT} is the regulation power consumed by wavelength tuning and ϵ is the tuning efficiency in mW/nm. Both of them are typically known in ONoCs. Therefore, with the obtained power loss before tuning (ΔP) for every MR, the thermal sensitivity model (Equation (4)), as well as our methodology, is still feasible.

To simplify this process, we further extend Equation (4) to Equation (10) with fine-grained consideration of the thermal sensitivity of MRs in the presence of local wavelength tuning,

$$T = T_0 + \frac{\sqrt{\left(\beta + \frac{P_{WT}^2}{\epsilon^2} \right) \cdot c^{(k \times \Delta P_{tuned})}} - \beta + (\lambda_{in} - \lambda_0)}{\rho_{MR}}. \quad (10)$$

For every MR, once we know the regulation power consumed for wavelength tuning technique (P_{WT}) and obtain the optical power loss of the MR after tuning (ΔP_{tuned}) using the CTS, the ambient temperature of the MR can be derived from Equation (10) directly.

7 CONCLUSION

In this article, we have proposed a novel hardware-software collaborative solution for run-time chip temperature monitoring. By utilizing the intrinsic thermal sensitivity of MRs and the inter-processor communications in ONoCs, the proposed collaborative thermal sensing technique achieves high monitoring accuracy and efficiency, while requires no additional hardware or chip area but lightweight software computations. Thanks to these favorable properties, our technique can also be used as an efficient complement of existing thermal sensing techniques. Experimental results based on professional photonic simulations strongly validate the effectiveness of the proposed technique. In future work, we plan to evaluate our technique based on fabricated optical devices and systems.

REFERENCES

- [1] [n.d.]. FDTD Solutions. Retrieved from <https://www.lumerical.com/products/fdtd/>.
- [2] [n.d.]. HEAT. Retrieved from <https://www.lumerical.com/products/heat/>.
- [3] [n.d.]. INTERCONNECT. Retrieved from <https://www.lumerical.com/products/interconnect/>.
- [4] [n.d.]. ANSYS®, Inc. Retrieved from <http://www.ansys.com/>.
- [5] Andre L. Aita, Michiel A. P. Pertijs, Kofi A. A. Makinwa, and Johan H. Huijsing. 2009. A CMOS smart temperature sensor with a batch-calibrated inaccuracy of $\pm 0.25^\circ\text{C}$ (3σ) from -70°C to 130°C . In *Proceedings of the International Solid-State Circuits Conference (ISSCC'09)*. IEEE, 342–343.
- [6] Tejasvi Anand, Kofi A. A. Makinwa, and Pavan Kumar Hanumolu. 2016. A VCO based highly digital temperature sensor With $0.034^\circ\text{C}/\text{mV}$ supply sensitivity. *IEEE J. Solid-State Circ.* 51, 11 (2016), 2651–2663.
- [7] Keren Bergman, Luca P. Carloni, Aleksandr Biberman, Johnnie Chan, and Gilbert Hendry. 2014. *Photonic Network-on-chip Design*. Springer.
- [8] Wim Bogaerts, Peter De Heyn, Thomas Van Vaerenbergh, Katrien De Vos, Shankar Kumar Selvaraja, Tom Claes, Pieter Dumon, Peter Bienstman, Dries Van Thourhout, and Roel Baets. 2012. Silicon microring resonators. *Laser Photon. Rev.* 6, 1 (2012), 47–73.
- [9] Poki Chen, Chun-Chi Chen, Yu-Han Peng, Kai-Ming Wang, and Yu-Shin Wang. 2010. A time-domain SAR smart temperature sensor with curvature compensation and a 3σ inaccuracy of $-0.4^\circ\text{C} + 0.6^\circ\text{C}$ over a 0°C to 90°C range. *IEEE J. Solid-State Circ.* 45, 3 (2010), 600–609.
- [10] Ching-Che Chung and Cheng-Ruei Yang. 2011. An autocalibrated all-digital temperature sensor for on-chip thermal monitoring. *IEEE Trans. Circ. Syst. II* 58, 2 (2011), 105–109.
- [11] G. Cocorullo and I. Rendina. 1992. Thermo-optical modulation at $1.5 \mu\text{m}$ in silicon etalon. *Electr. Lett.* 28, 1 (1992), 83–85.
- [12] Thomas H. Cormen. 2009. *Introduction to Algorithms*. MIT Press.
- [13] Huaxi Gu, Kwai Hung Mo, Jiang Xu, and Wei Zhang. 2009. A low-power low-cost optical router for optical networks-on-chip in multiprocessor systems-on-chip. In *Proceedings of the IEEE Computer Society Annual Symposium on VLSI*. 19–24.

- [14] Huaxi Gu, Jiang Xu, and Wei Zhang. 2009. A low-power fat tree-based optical network-on-chip for multiprocessor system-on-chip. In *Proceedings of the Design, Automation & Test in Europe Conference & Exhibition (DATE'09)*. 3–8.
- [15] Pengxing Guo, Weigang Hou, Liang Guo, Wei Sun, Chuang Liu, Hainan Bao, Luan Duong, and Weichen Liu. 2019. Fault-tolerant routing mechanism in 3D optical network-on-chip based on node reuse. *IEEE Trans. Parallel Distrib. Syst.* (2019). Early access.
- [16] Pengxing Guo, Weigang Hou, Lei Guo, Qiang Yang, Yifan Ge, and Hang Liang. 2018. Low insertion loss and non-blocking microring-based optical router for 3d optical network-on-chip. *IEEE Photon. J.* 10, 2 (2018), 1–10.
- [17] Wei Huang, Shougata Ghosh, Sivakumar Velusamy, Karthik Sankaranarayanan, Kevin Skadron, and Mircea R Stan. 2006. HotSpot: A compact thermal modeling methodology for early-stage VLSI design. *IEEE Trans. VLSI Syst.* 14, 5 (2006), 501–513.
- [18] Hyun-Tae Kim and Miao Yu. 2016. Cascaded ring resonator-based temperature sensor with simultaneously enhanced sensitivity and range. *Opt. Expr.* 24, 9 (2016), 9501–9510.
- [19] S. J. Koester, Laurent Schares, Clint Lee Schow, Gabriel Dehlinger, and R. A. John. 2006. Temperature-dependent analysis of Ge-on-SOI photodetectors and receivers. In *Proceedings of the International Conference on Group IV Photonics*. IEEE, 179–181.
- [20] Hasnain Lakdawala, Y. William Li, Arijit Raychowdhury, Greg Taylor, and Krishnamurthy Soumyanath. 2009. A 1.05 V 1.6 mW, 0.45°C 3 σ Resolution $\Sigma\Delta$ Based Temperature Sensor With Parasitic Resistance Compensation in 32 nm Digital CMOS Process. *IEEE J. Solid-State Circ.* 44, 12 (2009), 3621–3630.
- [21] Yu-Min Lee, Tsung-Heng Wu, Pei-Yu Huang, and Chi-Ping Yang. 2013. NUMANA: A hybrid numerical and analytical thermal simulator for 3-D ICs. In *Proceedings of the Design, Automation & Test in Europe Conference & Exhibition (DATE'13)*. IEEE, 1379–1384.
- [22] Mengquan Li, Weichen Liu, Lei Yang, Peng Chen, and Chao Chen. 2017. Chip temperature optimization for dark silicon many-core systems. *IEEE Trans. Comput.-Aid. Des. Integr. Circ. Syst.* 37, 5 (2017), 941–953.
- [23] Mengquan Li, Weichen Liu, Lei Yang, Peng Chen, Duo Liu, and Nan Guan. 2019. Routing in optical network-on-chip: Minimizing contention with guaranteed thermal reliability. In *Proceedings of the Asia and South Pacific Design Automation Conference (ASP-DAC'19)*. ACM, 364–369.
- [24] Sheng Li, Jung Ho Ahn, Richard D. Strong, Jay B. Brockman, Dean M. Tullsen, and Norman P. Jouppi. 2009. McPAT: An integrated power, area, and timing modeling framework for multicore and manycore architectures. In *Proceedings of the International Symposium on Microarchitecture*. ACM, 469–480.
- [25] Zhongqi Li, Amer Qouneh, Madhura Joshi, Wangyuan Zhang, Xin Fu, and Tao Li. 2015. Aurora: A cross-layer solution for thermally resilient photonic network-on-chip. *IEEE Trans. VLSI Syst.* 23, 1 (2015), 170–183.
- [26] Weichen Liu, Mengquan Li, Wanli Chang, Chunhua Xiao, Yiyuan Xie, Nan Guan, and Lei Jiang. 2019. Thermal sensing using micro-ring resonators in optical network-on-chip. In *Proceedings of the Design, Automation & Test in Europe Conference & Exhibition (DATE'19)*. IEEE, 1611–1614.
- [27] Weichen Liu, Lei Yang, Weiwen Jiang, Liang Feng, Nan Guan, Wei Zhang, and Nikil Dutt. 2018. Thermal-aware task mapping on dynamically reconfigurable network-on-chip based multiprocessor system-on-chip. *IEEE Trans. Comput.* 67, 12 (2018), 1818–1834.
- [28] Weichen Liu, Juan Yi, Mengquan Li, Peng Chen, and Lei Yang. 2018. Energy-efficient application mapping and scheduling for lifetime guaranteed MPSoCs. *IEEE Trans. Comput.-Aid. Des. Integr. Circ. Syst.* 38, 1 (2018), 1–14.
- [29] K. A. A. Makinwa. 2010. Smart temperature sensors in standard CMOS. *Proc. Eng.* 5 (2010), 930–939.
- [30] Moustafa Mohamed, Zheng Li, Xi Chen, Li Shang, Alan Mickelson, Manish Vachharajani, and Yihe Sun. 2010. Power-efficient variation-aware photonic on-chip network management. In *Proceedings of the International Symposium on Low-Power Electronics and Design (ISLPED'10)*. IEEE, 31–36.
- [31] Kishore Padmaraju and Keren Bergman. 2014. Resolving the thermal challenges for silicon microring resonator devices. *Nanophotonics* 3, 4–5 (2014), 269–281.
- [32] Santiago Pagani, Jian-Jia Chen, Muhammad Shafique, and Jörg Henkel. 2015. MatEx: Efficient transient and peak temperature computation for compact thermal models. In *Proceedings of the Design, Automation & Test in Europe Conference & Exhibition (DATE'15)*. IEEE, 1515–1520.
- [33] Assaf Shacham, Keren Bergman, and Luca P. Carloni. 2008. Photonic networks-on-chip for future generations of chip multiprocessors. *IEEE Trans. Comput.* 57, 9 (2008), 1246–1260.
- [34] Zhen Sheng, Liu Liu, Joost Brouckaert, Sailing He, and Dries Van Thourhout. 2010. InGaAs PIN photodetectors integrated on silicon-on-insulator waveguides. *Opt. Expr.* 18, 2 (2010), 1756–1761.
- [35] Uğur Sönmez, Fabio Sebastiano, and Kofi A. A. Makinwa. 2016. 11.4 1650 μm^2 thermal-diffusivity sensors with inaccuracies down to $\pm 0.75^\circ\text{C}$ in 40nm CMOS. In *Proceedings of the International Solid-State Circuits Conference (ISSCC'16)*. IEEE, 206–207.
- [36] Alexei Syrbu, A. Mereuta, V. Iakovlev, A. Caliman, P. Royo, and E. Kapon. 2008. 10 Gbps VCSELs with high single mode output in 1310 nm and 1550 nm wavelength bands. In *Proceedings of the Conference on Optical Fiber Communication/National Fiber Optic Engineers Conference (OFC/NFOEC'08)*. IEEE, 1–3.

- [37] Caspar P. L. van Vroonhoven, Dan d'Aquino, and Kofi A. A. Makinwa. 2010. A thermal-diffusivity-based temperature sensor with an untrimmed inaccuracy of $\pm 0.2^{\circ}\text{C}$ (3σ) from -55°C to 125°C . In *Proceedings of the International Solid-State Circuits Conference (ISSCC'10)*. IEEE, 314–315.
- [38] Yurii A. Vlasov and Sharee J. McNab. 2004. Losses in single-mode silicon-on-insulator strip waveguides and bends. *Opt. Expr.* 12, 8 (2004), 1622–1631.
- [39] Shipeng Wang, Xianglian Feng, Shiming Gao, Yaocheng Shi, Tingge Dai, Hui Yu, Hon-Ki Tsang, and Daoxin Dai. 2017. On-chip reconfigurable optical add-drop multiplexer for hybrid wavelength/mode-division-multiplexing systems. *Opt. Lett.* 42, 14 (2017), 2802–2805.
- [40] Shijun Xiao, Maroof H. Khan, Hao Shen, and Minghao Qi. 2007. Modeling and measurement of losses in silicon-on-insulator resonators and bends. *Opt. Expr.* 15, 17 (2007), 10553–10561.
- [41] Yiyuan Xie, Mahdi Nikdast, Jiang Xu, Wei Zhang, Qi Li, Xiaowen Wu, Yaoyao Ye, Xuan Wang, and Weichen Liu. 2010. Crosstalk noise and bit error rate analysis for optical network-on-chip. In *Proceedings of the Design Automation Conference (DAC'10)*. IEEE, 657–660.
- [42] Yiyuan Xie, Weihua Xu, Weilun Zhao, Yexiong Huang, Tingting Song, and Min Guo. 2015. Performance optimization and evaluation for torus-based optical networks-on-chip. *IEEE J. Lightw. Technol.* 33, 18 (2015), 3858–3865.
- [43] Yi Xu, Jun Yang, and Rami Melhem. 2012. Tolerating process variations in nanophotonic on-chip networks. In *Proceedings of the International Symposium on Computer Architecture (ISCA'12)*. IEEE, 142–152.
- [44] Lei Yang, Weichen Liu, Weiwen Jiang, Mengquan Li, Peng Chen, and Edwin Hsing-Mean Sha. 2016. FoToNoC: A folded torus-like network-on-chip based many-core systems-on-chip in the dark silicon era. *IEEE Trans. Parallel Distrib. Syst.* 28, 7 (2016), 1905–1918.
- [45] Yaoyao Ye, Jiang Xu, Xiaowen Wu, Wei Zhang, Xuan Wang, Mahdi Nikdast, Zhehui Wang, and Weichen Liu. 2012. System-level modeling and analysis of thermal effects in optical networks-on-chip. *IEEE Trans. Syst.* 21, 2 (2012), 292–305.
- [46] Chong Zhang, Shangjian Zhang, Jon D. Peters, and John E. Bowers. 2016. $8 \times 8 \times 40$ Gbps fully integrated silicon photonic network on chip. *Optica* 3, 7 (2016), 785–786.
- [47] Runjie Zhang, Mircea R. Stan, and Kevin Skadron. 2015. HotSpot 6.0: Validation, acceleration and extension. *University of Virginia, Tech. Report CS-2015-04* (2015).
- [48] Amirkoushyar Ziabari, Je-Hyoung Park, Ehsan K. Ardestani, Jose Renau, Sung-Mo Kang, and Ali Shakouri. 2014. Power blurring: Fast static and transient thermal analysis method for packaged integrated circuits and power devices. *IEEE Trans. VLSI Syst.* 22, 11 (2014), 2366–2379.

Received August 2019; revised September 2019; accepted September 2019