

# Parka: Thermally Insulated Nanophotonic Interconnects

Yigit Demir and Nikos Hardavellas

Northwestern University, Department of Electrical Engineering and Computer Science, Evanston, IL, USA  
yigit@u.northwestern.edu, nikos@northwestern.edu

**Abstract**—Silicon-photonics are emerging as the prime candidate technology for energy-efficient on-chip interconnects at future process nodes. However, current designs are primarily based on microrings, which are highly sensitive to temperature. As a result, current silicon-photonics interconnect designs expend a significant amount of energy heating the microrings to a designated narrow temperature range, only to have the majority of the thermal energy waste away and dissipate through the heat sink, and in the process of doing so heat up the logic layer, causing significant performance degradation to the cores and inducing thermal emergencies. We propose *Parka*, a nanophotonic interconnect that encases the photonic die in a thermal insulator that keeps its temperature stable with low energy expenditure, while minimizing the spatial and temporal thermal coupling between logic and silicon-photonics components. *Parka* reduces the microring energy by 3.8–5.4x and achieves 11–23% speedup on average (34% max) depending on the cooling solution used.

## I. INTRODUCTION

Silicon photonics have emerged as a promising solution to meet the growing demand for high-bandwidth, low-latency, and energy-efficient communication in manycore processors. Silicon waveguides are more efficient for long-distance on-chip communication than electrical signaling [24], and nanophotonic devices can be manufactured by simply adding a few new steps in the CMOS manufacturing process [5]. While silicon-photonics devices can be manufactured alongside CMOS logic even on the same die [5], designers typically assume a simplified process where the photonic components are housed within a photonic die, which is 3D-stacked to a logic die that contains cores, caches, and other electronic components.

Due to this arrangement, the thermal variations of the logic die directly couple to the photonic devices. These thermal variations may occur rapidly depending on the workload, are both spatial and temporal in nature, and can exceed 30 °C difference. Moreover, both processor and memory chips are susceptible to thermal variations [17]. As current silicon-photonics designs are predominantly based on microring resonators, these thermal fluctuations may prevent the optical interconnect from

functioning. Microrings are tuned to resonate at a particular wavelength when they are at a set temperature, but they are highly thermally sensitive devices. For example, the resonant wavelength of a microring modulator with 5  $\mu\text{m}$  diameter shifts by 0.11  $\text{nm}/^\circ\text{C}$  [18]. As a typical photodetector requires optical power no less than 3  $\text{dB}$  below the peak to operate properly, the microrings can withstand no more than 2.8 °C of temperature shift assuming 5  $\text{nm}$  wavelength separation.

To keep the microrings resonating at their appropriate wavelengths designers employ trimming, a technique that dynamically shifts the microring's resonant wavelength towards the red through heating, or shifts it towards the blue through current injection. Trimming by current injection causes instability and thermal runaways [20], thus microrings are typically maintained at a constant temperature using the heaters only. Modulators with integrated heaters have been shown to produce error-free 10  $\text{Gb/s}$  modulation across a 60 °K temperature variation range, with comparable tuning efficiencies [38]. Because only the heaters are used, the microrings are tuned to temperatures above the maximum temperature that the microprocessor reaches.

Unfortunately this means that the heaters need to work continuously to keep the microrings at such high temperature, and at the same time the majority of the heating power is wasted as it dissipates through the package to the heat sink. As a result, it is common for microring heaters to consume upwards of 40  $\text{W}$  [20], the majority of which is wasted. To make matters worse, this thermal energy heats up the logic layer to temperatures very close to its operational limit, which forces the system to throttle the cores, thereby reducing performance. The runaway heat also increases the frequency and magnitude of thermal emergencies, and accelerates the aging of the logic die.

The solution we propose is rather simple: thermally decouple the 3D-stacked logic die from the photonics die by introducing an insulating layer between them to maintain higher thermal stability and easier trimming. More specifically, our contributions are:

- We propose *Parka*, a nanophotonic NoC that encases the photonic die in a thermal insulator that keeps its temperature stable with low energy expenditure, while minimizing the spatial and temporal thermal coupling between logic and silicon-photonics components.
- We quantify the ring heating power consumption for a large-scale multicore under a variety of insulation methods and cooling solutions.
- We evaluate the performance impact of thermal decoupling on a multicore running a range of scientific workloads, under realistic physical constraints.

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 the author(s) 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](mailto:Permissions@acm.org).

NOCS '15, September 28 - 30, 2015, Vancouver, BC, Canada  
Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-3396-2/15/09 \$15.00

DOI: <http://dx.doi.org/10.1145/2786572.2786597>




Fig. 1. Proposed Parka Architecture.

Our results indicate that Parka reduces the ring heating power by 3.8-5.4x on average across our workload suite. Moreover, the energy savings allow for providing a higher power budget to the cores, which enables them to run faster. Parka on a radix-16 crossbar allows the multicore to achieve 11-23% speedup (34% max) over a baseline scheme with no insulation, depending on the cooling solution used.

## II. PHOTONIC DIE INSULATION WITH PARKA

The basic building block of silicon-photonics interconnects is the microring resonators, which are designed to resonate at a specific wavelength to realize add/drop filters and modulators. The microring resonators are very susceptible to temperature changes, because the refractive index of Si changes with temperature, in turn changing the resonance wavelength. Trimming keeps the microrings resonating at their appropriate wavelengths by dynamically shifting the microring’s resonant wavelength towards the red through heating, or towards the blue through current injection. Microrings are typically kept at a constant temperature using the heaters only, as current injection causes instability and thermal runaways [20]. The strong thermal coupling of the logic and photonic dies means that trimming by heating requires that the photonic die is heated to a temperature above the maximum temperature of the logic die.

This ring-heating power is mainly wasted, as it dissipates through the processor stack into the logic layer and eventually through the heat-sink, which is designed to remove heat from the processor stack. This heats the logic layer close to the limits of safe operating temperatures. A thermal emergency occurs when the logic die temperature exceeds the safe operation limits, at which point the cores are throttled or turned off to lower the temperature. Therefore, high ring heating power consumption makes the multicore processor more susceptible to thermal emergencies, and may decrease its performance significantly.

Parka reduces the wasted energy and the heating of the logic layer by thermally decoupling the 3D-stacked logic die from the photonics die by placing an insulation layer between them (Figure 1). The insulation layer increases the thermal resistivity of the heat path from the photonics layer to the heat sink, and (a) allows for easier microring trimming by trapping the heat within the photonics layer, (b) reduces the temperature variation in the photonics layer, and (c) minimizes the heating of the logic die induced by the microring heaters. The processor die is placed close to the heat sink to allow better cooling, while an oxidized macro porous Si layer [19] realizes the ther-

mal insulation, as porous Si has 100x lower thermal conductivity than Si [19]. The porous Si layer is 150  $\mu\text{m}$  thick, as we find that a thinner layer does not provide adequate thermal insulation. The power delivery and communication between the dies is maintained through high aspect ratio TSVs [12, 28, 35].

Adding the insulation layer is expected to increase the manufacturing cost only marginally. The porous Si insulation layer can be readily integrated into the CMOS process by passing a plain silicon die through a simple electrochemical process that oxidizes it [19]. This silicon die is not subject to the regular yield-induced costs of dies that implement complex logic and require multiple mask exposures and several metal layers, and thus it is significantly cheaper. The addition of the porous Si layer also does not affect the number of TSVs and the number of pins in the package, which together with the logic and photonic dies constitute the dominant cost factors [10,36]. The thickness of the insulation layer impacts the TSVs’ height, but the cost is highly insensitive to it [10,36]. The additional layer will incur 3D-bonding costs, but these will increase the total cost by less than 1.5% [10,36].

Insulation can be achieved also by a 5  $\mu\text{m}$ -thick air or vacuum cavity etched between layers, a technique for which prototypes have been successfully manufactured and characterized [35]. Air has a thermal resistivity of 40  $\text{m-K/W}$ , which is 40 times higher than porous Si, so it would be an even better insulator. However, this technique is more challenging to employ than oxidized porous Si. Thus, we maintain our conservative assumptions using porous Si insulators and do not consider alternative insulation techniques further. It is important to note that Parka does not depend on the exact insulator technology used. As processes mature and better materials and techniques become available, they can be employed by Parka to achieve even higher power savings than the ones we show in this paper.

## III. EXPERIMENTAL METHODOLOGY

### A. Ring-Heater Power Consumption Analysis

We model a photonic die with microrings tuned to 90  $^{\circ}\text{C}$  (363.15  $^{\circ}\text{K}$ ), which is the maximum temperature that the logic die can reach. To calculate the total ring heating power we extend the method by Nitta *et al.* [20] by estimating the ring-heater power consumption while accounting for the heating of the photonic die by the operation of the cores. While one can assume that the heaters are employed to shift the resonant wavelengths of the microrings only momentarily according to the local temperature, keeping a stable temperature for the die as a whole is a more realistic approach [20].

We model a multicore where 50  $\mu\text{m}$ -thick logic and photonic dies are 3D-stacked, and separated by a 150  $\mu\text{m}$  porous Si insulation layer, as shown in Figure 1. The thermal resistivity is 0.01  $\text{m-K/W}$  for Si, and 1  $\text{m-K/W}$  for the porous Si insulator [19]. We evaluate the ring-heater power consumption of Parka using the 3D extension of HotSpot [29], a thermal modeling tool based on an equivalent circuit of thermal resistances and capacitances. We evaluate Parka’s impact on the heat transfer rate between dies via a transient thermal analysis at 300  $\mu\text{s}$  time steps. The ambient temperature is fixed at 45  $^{\circ}\text{C}$  (318.15  $^{\circ}\text{K}$ ).

TABLE 1. ARCHITECTURAL PARAMETERS.

<b>CMP Size</b>	64 cores, 480mm <sup>2</sup>
<b>Processing Cores</b>	ULTRASPARC III ISA, up to 5Ghz, OoO, 4-wide dispatch/retirement, 96-entry ROB
<b>L1 Cache</b>	Split I/D, 64KB 2-way, 2-cycle load-to-use, 2 ports, 64-byte blocks, 32 MSHRs, 16-entry victim cache
<b>L2 Cache</b>	Shared, 512 KB per core, 16 way, 64-byte blocks, 14 cycle-hit, 32 MSHRs, 16-entry victim cache
<b>Memory Controllers</b>	One MC per 4 cores, uniformly distributed, 1 channel per MC, round-robin page interleaving
<b>Main Memory</b>	Optically connected memory [1], 10 ns access
<b>Networks</b>	SWMR crossbar, radix-16

Our model accounts for the thermal impact of TSVs, as they are highly conductive, and also for the individual ring trimming power required to overcome process variations, as described in [14]. We model a design that employs a total of 76,800 microrings, which are driven by one TSV each. We model high-aspect ratio TSVs with 10  $\mu\text{m}$  diameter [28]. All the TSVs together cover a 6  $\text{mm}^2$  area, which corresponds to 1.25% of the chip area and contributes only 0.5% to the total cost [10,36]. It is important to note that this is not an overhead that Parka imposes to the system; rather, it is the overhead of 3D-stacking the photonic and the logic dies, and it is incurred by both Parka and the baseline system.

### B. Multicore System Performance and Energy Analysis

To evaluate the impact of Parka on a realistic multicore system, we model a multicore processor on a full-system cycle-accurate simulator based on Flexus 4.0 [13, 33] integrated with Booksim 2.0 [6] and DRAMSim 2.0 [25]. Figure 2 describes our simulation tool chain. We target a 16  $\text{nm}$  technology, and have updated our tool chain accordingly based on ITRS projections [11]. We collect runtime statistics from full-system simulations, and use them to calculate the power consumption of the system using McPAT [16], and the power consumption of the optical networks using the analytical power model by Joshi *et al.* [14]. The analytical model we use for the power calculation of the photonic components results in similar overall power estimates as DSENT [30], but it also provides an easy breakdown of the power consumed by each one of the nanophotonic components in our network. We estimate the temperature of the chip using the 3D extension of HotSpot 5.0 [29]. The estimated

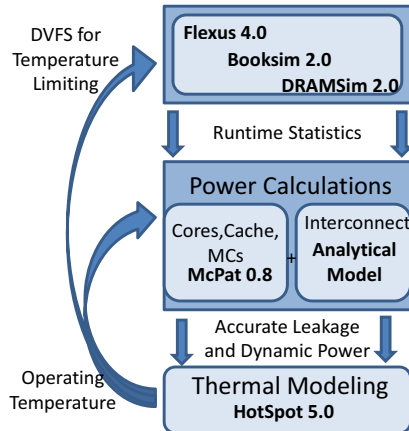


Fig. 2. Simulation Flow Chart.

TABLE 2. WORKLOAD DETAILS.

Suite	Workload	Description
NAS	appbt	Independent equations system solver 32x32x32 grid, 1e-12 tolerance, 8e-4 time step, 1.2 SSOR iteration relaxation factor
SPEC-CPU	tomcatv	Vectorized mesh generation; parallel version of 101.tomcatv from SPEC-FP 4,096 array size, 10 iterations
SPLASH-2	barnes	Barnes-Hut hierarchical N-body simulation 64K particles., 2.0 subdiv. tol., 10.0 fleaves, 2.0 fcells, 0.025 time step, 0.05 softening
	fmm	Particle simulation via adaptive fast multipole 131K particles, two clusters, plummer distr., 1e-6 precision, 30 steps, 0.025 step duration
	ocean	Eddy & boundary oceanic currents simulator 1026 x 1026 grid, 20,000 meters, 9,600 sec, 1e-7 tolerance
PARSEC	bodytrack	Annealed particle filter to track human body 4 cameras, 4 frames, 4,000 particles, 5 annealing layers (simlarge)
Other Scientific	moldyn	Molecular dynamics simulation 19,652 molecules, max interactions 3,200,000
	em3d	Electromagnetic force simulation 768K nodes, degree 2, span 5, 15% remote

temperature is then used to refine the leakage power estimate. We adjust the voltage and frequency of the logic die based on the stable-state power and temperature estimates (Figure 2), and we repeat the process until the system reaches a stable state and additional iterations result in no further changes on temperature and overall power consumption.

Using the methodology above, we simulate a 64-core multicore system. By scaling existing core designs down to 16  $\text{nm}$  we estimate that 64 cores would require a 480  $\text{mm}^2$  die. Table 1 details the architectural modeling parameters. We model realistic multicore systems that employ dynamic thermal management by throttling the voltage and the frequency of the chip to keep it within safe operational temperatures (below 90  $^{\circ}\text{C}$ , i.e., 363.15  $^{\circ}\text{K}$ ). The simulated multicore executes a selection of SPLASH-2 and PARSEC benchmarks, and other scientific workloads. The workload parameters are detailed in Table 2.

### C. Interconnect and Nanophotonic Parameters

We employ a cycle-accurate network simulator based on Booksim 2.0 [6], which models a radix-16 SWMR crossbar. The simulator models a single-cycle router, with 1-cycle E/O and O/E conversions. We assume a 480  $\text{mm}^2$  chip, which employs a 10  $\text{cm}$  waveguide with a round trip time of 5 cycles. The link latency (1-5 cycles) is calculated based on the tra-

TABLE 3. NANOPHOTONIC PARAMETERS.

	per Unit	Radix-16 Total
<b>DWDM</b>		16
<b>WG Loss</b>	0.3 dB/cm[3]	3 dB
<b>Nonlinearity</b>	1 dB	1 dB
<b>Modulator Ins.</b>	0.5 dB	0.5 dB
<b>Ring Through</b>	0.01 dB	2.56 dB
<b>Filter Drop</b>	1.2 dB	1.2 dB
<b>Photodetector</b>	0.1 dB	0.1 dB
<b>Total Loss</b>		8.36 dB
<b>Detector</b>		-20 dBm
<b>Mod./Demod. Energy (10 GHz)</b>		150 fJ/bit

versed waveguide length. The buffers are 20-flits deep, with a flit size of 300 bits. The maximum core frequency is 5 GHz, and the optical interconnect runs at 10 GHz. We derive the nanophotonic parameters from [1] and detail them in Table 3. The data bus is 300-bits wide (300 wavelengths with 16-way DWDM) powered by an off-chip laser source.

Unfortunately, there is little consensus on the optical loss parameters used or projected in literature, as parameters exhibit a variance over 10x across publications. However, the design of an optical interconnect highly depends on the losses of the optical components used. If the off-ring through loss on the radix-16 crossbar was 10x higher (i.e., 0.1dB), the interconnect wouldn't employ 64-way DWDM, as this would increase the laser power to unsustainable levels. Rather, it would be optimized with a lower DWDM (using more waveguides), keeping the total optical loss (and hence laser power) the same. In our work we limit the network to 16 DWDM because the number of turned-off rings on a single optical path of a crossbar is high, so limiting the DWDM helps keep the total optical loss at reasonable levels. 16-way DWDM has already been demonstrated and it is a widely-accepted parameter.

#### D. Modeling Cooling Solutions

The ring-heating power requirement depends highly on the cooling solution. Aggressive cooling solutions are capable of faster heat removal from the processor stack, which is likely to force the ring heaters to work even harder to keep the photonic layer at the tuned temperature. Therefore, the thermal decoupling that Parka advocates will be more important when better cooling solutions are employed. To evaluate the impact of Parka across cooling solutions we model both forced-air cooling (convective thermal resistance  $R_{\text{conv}} = 0.25 \text{ K/W}$ ) and a liquid cooling solution ( $R_{\text{conv}} = 0.15 \text{ K/W}$  [27]).

For the liquid cooling solution we assume that microchannels facilitate forced convective interlayer cooling with single-phase fluids, in particular water. While other single-phase fluids with higher thermal capacitance exist, they are toxic and thus impractical to deploy. We model high-aspect ratio TSVs with 10  $\mu\text{m}$  diameter [28], located and etched within 100  $\mu\text{m}$ -wide microchannel walls as in [26]. We assume uniformly distributed microchannels, and equivalent fluid flow rate through each channel in the same layer. Although variation of the fluid flow due to nonuniform heat flux can exist, variations stay below 2% for single-phase flows and have negligible impact on the cooling system's performance [26]. The fluid pump and valve consume 1.3 W per 10 ml/min flow, and the power is linear to the volumetric fluid flow [26].

### IV. EXPERIMENTAL RESULTS

#### A. Impact on the Ring-Heating Power Consumption

Parka thermally decouples the photonics die from the processor die using a porous Si insulating layer which reduces the thermal fluctuations caused by the processor layer, and traps the heat in the photonics die allowing for easier trimming. In this section we evaluate the ring-heating power consumption of Parka on a 64-core processor, and compare it against an architecture with no insulation.




Fig. 3. Transient analysis of temperature fluctuations in the photonics die.

First we evaluate the thermal shielding effect of the insulating layer by observing the temperature variation in the photonics die resulting from temperature fluctuations in the processor die. We increase the power consumption in the processor layer (from its idle level) to its maximum allowed level, and observe the temperature change in the photonics layer (Figure 3). The processor die stays at 66 °C (339.15 °K) when in the idle state, and its temperature reaches 90 °C (363.15 °K) rapidly when it is turned on (~18 ms). The temperature of the photonics die closely tracks the temperature change of the processor die when there is no insulation. However, for Parka, it takes twice as long for photonics layer to reach 90 °C (Figure 3), because of the thermal shielding effect of the insulating layer. Note that the insulating layer not only shields the fluctuations towards the higher temperature levels, but it also shields from the dips in the temperature. Overall, Parka allows for easier trimming because it shields the photonics layer from the short temperature fluctuations occurring in the logic silicon layer.

Thermally decoupling the photonics layer from the rest of the processor stack allows for trimming with less ring heater power consumption, because it does not allow the heat generated by the ring heaters dissipate through the heat sink easily. The insulating layer increases the thermal resistance on the heat path to the heat sink, so it traps the heat within the photonics die. Therefore, Parka's ring-heaters can bring the whole photonics die to a stable temperature level which is higher than the maximum execution temperature at the processor layer with less power. Figure 4 shows a scenario where we present both the shielding and heat trapping effect of Parka. Figure 4.a shows a snapshot (at time  $t_0$ ) of the thermal map of the processor die when running a real workload (appbt). We assume that at time  $t_0$  all processors stop, and they only dissipate leakage power until time  $t_1$ . We estimate that the processor die leakage power is ~30 W when idle. Figure 4.b shows the temperature maps of the photonics layer at time  $t_1$ . We observe that the photonics layer stays at a higher temperature for Parka compared to no insulation, as it retains the heat due to the insulating layer.

In the example in this figure we assume that the ring heaters are also off until time  $t_1$ . At time  $t_1$ , the ring heaters are turned on to bring the photonics layer to a stable 90 °C (363.15 °K), and Figure 4.c shows the power distribution of these ring heaters. We observe that Parka requires less ring-heating power. There are two reasons for this: first, the photonics layer is at a

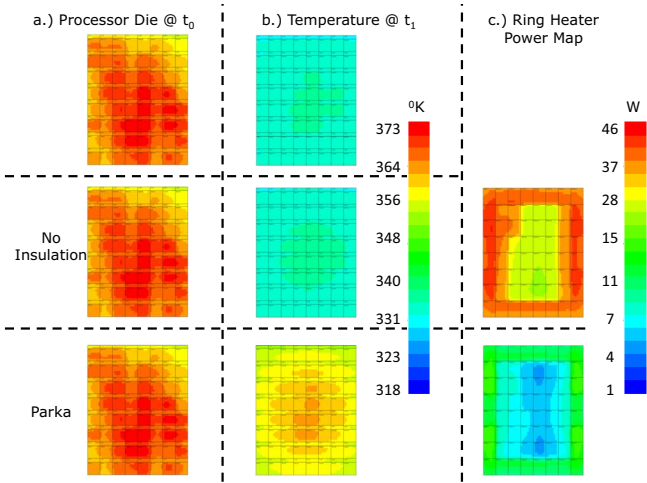


Fig. 4. Case study: Impact of thermal insulation on the photonics layer temperature and the ring-heating power consumption.

higher temperature at time  $t_1$ , so there is a smaller temperature difference (to  $90^\circ\text{C}$ ) to cover. Second, it is easier to close this temperature difference with Parka because the heat generated by the ring heaters stays within the photonics die.

The amount of ring-heating power required to keep the photonics layer at a stable  $90^\circ\text{C}$  highly depends on the power consumption of the processor die. When the processor die is idle, the ring heaters have to work harder to warm up the photonics die. In Figure 5, we show the ring-heating power for a range of power consumption levels of the processor die. We observe that for every processor die utilization level, Parka consumes less ring-heating power than the no-insulation case. The maximum amount of ring-heating power required for Parka is 3.5x lower than the maximum ring-heating power required without insulation.

It is important to note that the ring-heating power requirement highly depends on the cooling solution. Better cooling solutions are capable of removing the heat from the processor stack at higher rates, which may force ring heaters to work even harder to keep the photonics layer warm. Therefore, the thermal decoupling will be more important when better cooling solutions such as liquid cooling are employed. To observe this effect we repeat the same ring-heating power estimation with a liquid cooling solution (liquid cooling  $R_{\text{conv}} = 0.15 \text{ K/W}$  [27], whereas forced-air cooling  $R_{\text{conv}} = 0.25 \text{ K/W}$ ).

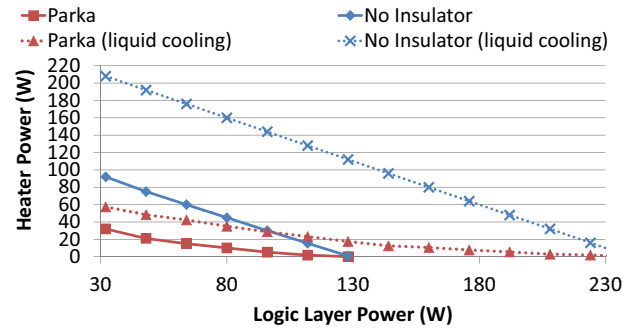



Fig. 5. Ring-Heating Power vs. Processor Die Power

We observe that with liquid cooling the operational temperature at the processor layer stays under  $90^\circ\text{C}$  when the processor die consumes up to  $250 \text{ W}$  (Figure 6.a), while forced-air cooling can sustain at best only up to  $130 \text{ W}$  and passive cooling less than  $100 \text{ W}$ . More importantly, we observe that the magnitude of the thermal fluctuations on the processor layer is higher under an aggressive cooling solution, because higher utilization levels are permitted within the power budget, and the idle temperature is lower due to better cooling. Figure 6.b shows the instructions per second attained during the execution of a given code fragment of an application (appbt) when liquid or forced-air cooling are employed. We observe that liquid cooling allows higher performance, but also that the temperature under liquid cooling fluctuates between  $54\text{--}90^\circ\text{C}$ , while for the same exact execution segment run under forced-air cooling the temperature fluctuates between  $68\text{--}90^\circ\text{C}$ . Thus, the temperature fluctuation range on the simulated multicore is  $14^\circ\text{C}$  wider with liquid cooling compared to forced-air cooling when running the same code fragment (Figure 6.b).

As processor temperatures fluctuate during execution, the ring-heaters have to step in to keep the photonics layer at a stable temperature. We analyze this effect by running a collection of diverse workloads on our simulated multicore system and calculating the average ring-heating power consumed by each application (Figure 7). We observe that the temperature fluctuations are higher when running memory-intensive workloads (e.g., bodytrack, em3d, ocean, appbt), hence the ring-heating power consumption is also higher. On average ring heaters consume  $16.9 \text{ W}$  ( $22.4 \text{ W}$  maximum) when there is no insulation. Parka allows for easier trimming by shielding from short fluctuations and trapping the heat, so it consumes on average 3.8x less ring heating power ( $4.4 \text{ W}$  on average).

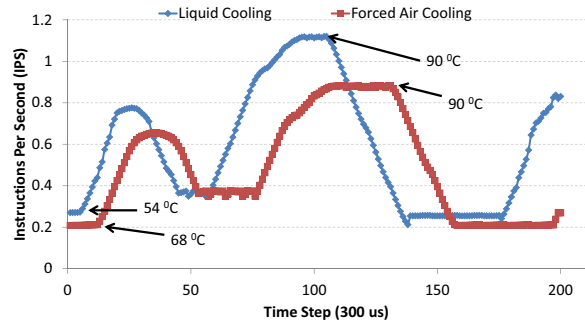


Fig. 6. (a) Processor die temperature vs. power consumption, (b) Instructions per second trace and thermal variation for the same execution segment (appbt).




Fig. 7. Average ring-heating power consumption of real-world applications

The liquid cooling solution keeps the processor cooler and allows for cores to run faster, however this results in higher temperature fluctuations at the photonics layer. On top of that, with better heat dissipation from the photonics layer, ring-heaters have to consume more power to keep the photonics layer at a stable temperature. Figure 7 shows that the ring heaters have to consume  $28.2\text{ W}$  on average when there is no insulation and a liquid cooling solution is employed. However, employing an insulating layer in this case reduces the ring-heating power consumption by 5.4x on average ( $5.2\text{ W}$ ). Thus, Parka is essential when using aggressive cooling solutions.

### B. Impact on the Processor Temperature

Ring heaters warm up and keep the photonics die at a slightly higher temperature than the maximum operating temperature of the processor [20]. However, while heating the photonics die, the ring heaters also heat the processor die when there is no insulation. Heating the processor die forces it to operate close to its maximum operating temperature, even when it is idle. In this case, even a small increase in the utilization can cause a temperature spike which pushes the processor out of the safe operating limits causing it to throttle, and reducing performance. Thus, in the absence of an insulating layer the processor becomes highly vulnerable to thermal emergencies.

On the other hand, ring heaters consume 3.8x less power on average with Parka, and thus the processor layer remains cooler, because the overall power consumption in the processor stack is lower (leakage power is exponentially related to temperature). For example, Figure 8 shows that when compute components consume  $90\text{ W}$  at the logic layer, the ring heaters consume  $36\text{ W}$  when there is no insulation, but only  $7.2\text{ W}$  with

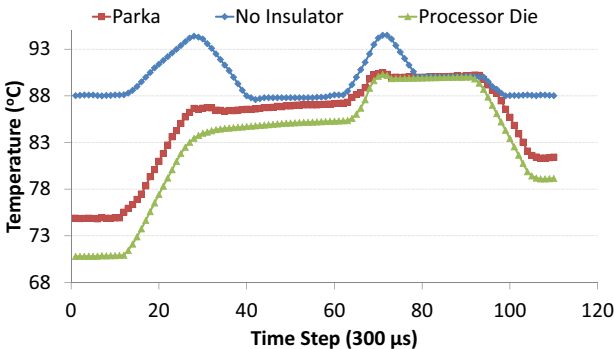


Fig. 9. Temperature trace (appbt) presenting thermal emergencies in a multicore, and the percentage of execution time spent under thermal emergencies.

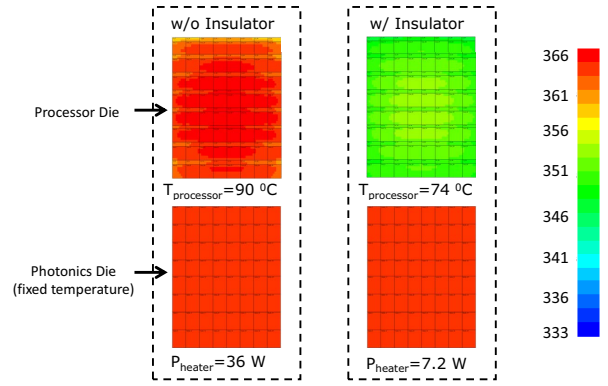


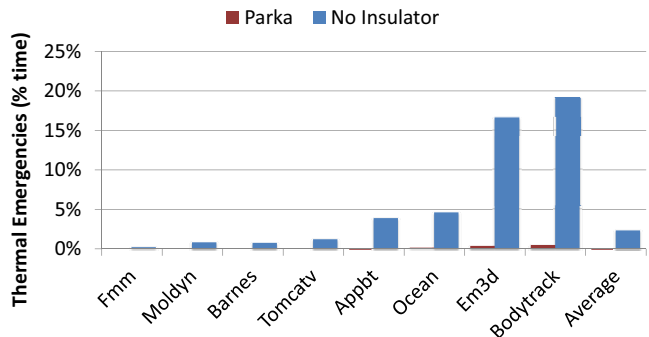
Fig. 8. Parka’s impact on the processor die temperature.

Parka. As a result, the logic layer stays at  $74\text{ }^{\circ}\text{C}$  ( $347.15\text{ }^{\circ}\text{K}$ ) with Parka, while it reaches  $\sim 90\text{ }^{\circ}\text{C}$  without insulation.

The ring heaters keep the processor die very close to the limit of safe operating temperature, so any increase in the processor utilization can push the processor into thermal emergencies. We present such an example in the execution window shown in Figure 9.a. The activity increase around time steps 12 and 62 push the processor temperature over  $90\text{ }^{\circ}\text{C}$  when there is no insulation, whereas with Parka the processor stays cooler and avoids the thermal emergencies. When running real applications, the processor runs into thermal emergencies up to 19% of the execution time (2% on average) when there is no insulation (Figure 9.b). The cores need to be throttled or completely turned off during a thermal emergency to allow for the processor to cool down and avoid permanent damage, so we expect that these thermal emergencies will significantly reduce the processor’s performance. In contrast, Parka’s processor die largely avoids thermal emergencies, and only experiences them for less than 1% of the execution time (Figure 9.b).

### C. Impact on a Realistic Multicore

Under realistic thermal and power constraints, the dynamic thermal management system in the processor throttles the cores to keep the chip within a safe temperature. The insulating layer, however, reduces the ring-heating power and results in a cooler chip, causes less core throttling, and provides higher performance. Overall, Parka reduces the ring-heating power consumption by 3.8x, which allows for the cores to run faster. As a result, the processor with the insulating layer runs 11% faster






Fig. 10. Realistic Multicore Performance with Parka.

on average (18% maximum) than the processor without the insulating layer (Figure 10). The ring-heating power consumption is higher when an aggressive cooling solution (e.g., liquid cooling) is employed, so the power savings of Parka are also higher. With liquid cooling, Parka outperforms the processor without the insulation by 23% on average (34% maximum).

## V. RELATED WORK

Silicon photonics are emerging as a promising technology for high-bandwidth, low-latency, and energy-efficient communication in multicore processors. Many different topologies that have been proposed, such as Corona [32] and many others [31,23,22], implement a nanophotonic MWSR crossbar topology for on-chip communication. Firefly [24] uses partitioned SWMR optical crossbars to connect clusters of electrically-connected mesh networks. Batten *et al.* [1,2] connect a many-core processor to DRAM memory using SWMR crossbars.

The high laser and ring-heating power consumption reduce the energy efficiency of the nanophotonic interconnects. Zhou *et al.* [37] identify the constant laser power consumption as an inefficiency, and propose a mechanism to increase average channel utilization by controlling active splitters to tune bandwidth on a binary tree network. Kurian *et al.* [15] propose an optical SWMR crossbar and electrical hybrid network, and mention that a Ge-based laser can be controlled to improve the laser energy efficiency. Joshi *et al.* [4] propose a scheme to distribute laser power across multiple busses based on the utilization levels to provide higher bandwidth and achieve higher energy efficiency. Zhang *et al.* [34] investigate the temperature gradients for the multicores with photonic interconnects, and propose a temperature-aware job allocation scheme to minimize the temperature gradients among the ring resonators. Demir and Hardavellas [7,9,8] advocate laser gating as an effective technique to eliminate the energy waste of laser sources and improve the energy efficiency of optical interconnects. All these works are orthogonal to Parka and can be used in addition to Parka to achieve even higher energy efficiency. Parka covers the photonic die with an insulation layer that keeps its temperature stable with low energy expenditure, while minimizing the spatial and temporal thermal coupling between logic and silicon-photonic components.

There are several techniques that can be used to resolve the thermal challenges of the silicon microring resonator devices. Methods to reduce the thermal dependence of microrings to tolerable levels include athermalization using negative thermo-optic materials or the embedment of the microring in a thermally-balanced interferometric structure. However, it is challenging to integrate the necessary polymer and TiO<sub>2</sub> materials into a CMOS-compatible fabrication process, and the interferometric structure still suffers from susceptibility to fabrication tolerances, increases the footprint of the microring, and it is challenging to adapt the technique to larger microring switch fabrics [21]. Thus, control-based techniques that aim to detect and react to the resonance shift due to thermal fluctuations are preferable, and several prototypes have been shown to withstand thermal variations across a wide temperature range up to 32–60 °K [38]. It is beyond the scope of this paper to provide a detailed review and comparison of such techniques. However, the interested reader could refer to some of the excellent surveys on this topic that are available in the literature, e.g., Padmaraju and Bergman [21].

## VI. CONCLUSION

Silicon-photonics are rapidly becoming a serious contender for energy-efficient on-chip interconnects at future process nodes. However, current designs are primarily based on microrings, which are highly sensitive to temperature. As a result, current silicon-photonic interconnect designs expend a significant amount of energy heating the microrings to a designated narrow temperature range, only to have the majority of the thermal energy waste away and dissipate through the heat sink, and in the process of doing so heat up the logic layer, causing significant performance degradation to the cores and inducing thermal emergencies. To address this problem we propose Parka, a nanophotonic NoC that encases the photonic die in a thermal insulator that keeps its temperature stable with low energy expenditure, while minimizing the spatial and temporal thermal coupling between logic and silicon-photonic components. Our results indicate that Parka reduces the ring heating power by 3.8-5.4x on average across our workload suite, depending on the cooling solution used, thereby eliminating a significant source of power waste in silicon-photonic interconnects. Moreover, by eliminating the microring-induced heating of the logic layer, Parka allows for providing a higher power budget to the cores, which enables them to run faster. Parka on a radix-16 crossbar allows the multicore to achieve 11-23% speedup (34% max) over a baseline scheme with no insulation, depending on the cooling solution used. We also observe that as cooling solutions become more aggressive in future technologies, the impact and importance of Parka increases.

## VII. ACKNOWLEDGEMENTS

This work was generously funded by NSF CAREER award CCF-1453853.

## REFERENCES

- [1] C. Batten, A. Joshi, J. Orcutt, A. Khilo, B. Moss, C. W. Holzwarth, M. A. Popovic, H. Li, H. I. Smith, J. L. Hoyt, F. X. Kartner, R. J. Ram, V. Stojanovic, and K. Asanovic. Building many-core processor-to-dram networks with monolithic cmos silicon photonics. *IEEE Micro*, 29(4):8-21, 2009.

- [2] C. Batten, A. Joshi, V. Stojanovic, and K. Asanovic. Designing chip-level nanophotonic interconnection networks. *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, 2(2):137- 153, 2012.
- [3] J. Cardenas, C. Poitras, J. Robinson, K. Preston, L. Chen, and M. Lipson. Low loss etchless silicon photonic waveguides. *Optics Express*, 17(6):4752- 4757, 2009.
- [4] C. Chen and A. Joshi. Runtime management of laser power in silicon-photonic multibus noc architecture. *IEEE Journal of Selected Topics in Quantum Electronics*, 19(2):3700713- 3700713, March 2013.
- [5] G. Chen, H. Chen, M. Haurylau, N. Nelson, P. M. Fauchet, E. Friedman, and D. Albonesi. Predictions of cmos compatible on-chip optical interconnect. In *7th International Workshop on System-Level Interconnect Prediction (SLIP)*, pages 13- 20, San Francisco, CA, 2005.
- [6] W. J. Dally and T. B. Principles and Practices of Interconnection Networks. Morgan Kaufmann Publishing Inc., 2004.
- [7] Y. Demir and N. Hardavellas. Ecolaser: An adaptive laser control for energy efficient on-chip photonic interconnects. In *Proceedings of the International Symposium on Low-Power Electronics and Design*, Aug. 2014.
- [8] Y. Demir and N. Hardavellas. Lac: Integrating laser control in a photonic interconnect. In *IEEE Photonics Conference (IPC)*, pages 28- 29, 2014.
- [9] Y. Demir and N. Hardavellas. Towards energy-efficient photonic interconnects. In *Proceedings of Optical Interconnects XV, SPIE Photonics West*, February 2015.
- [10] X. Dong, J. Zhao, and Y. Xie. Fabrication cost analysis and cost-aware design space exploration for 3-D ICs. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 29(12), December 2010.
- [11] European Semiconductor Industry Association (ESIA), Japan Electronics and Information Technology Industries Association (JEITA), Korean Semiconductor Industry Association (KSIA), Taiwan Semiconductor Industry Association (TSIA), and United States Semiconductor Industry Association (SIA). The international technology roadmap for semiconductors (itrs). <http://www.itrs.net/>, 2012 Edition.
- [12] A. C. Fischer, S. J. Bleiker, T. Haraldsson, N. Roxhed, G. Stemme, and F. Niklaus. Very high aspect ratio through-silicon vias (tsvs) fabricated using automated magnetic assembly of nickel wires. *Journal of Micromechanics and Microengineering*, 22(10):105001, 2012.
- [13] N. Hardavellas, S. Somogyi, T. F. Wenisch, R. E. Wunderlich, S. Chen, J. Kim, B. Falsafi, J. C. Hoe, and A. G. Nowatzky. SimFlex: a fast, accurate, flexible full-system simulation framework for performance evaluation of server architecture. *SIGMETRICS Performance Evaluation Review, Special Issue on Tools for Computer Architecture Research*, 31(4):31- 35, April 2004.
- [14] A. Joshi, C. Batten, Y.-J. Kwon, S. Beamer, I. Shamim, K. Asanovic, and V. Stojanovic. Silicon-photonic cmos networks for global on-chip communication. In *Proceedings of the IEEE International Symposium on Networks-on-Chip (NOCS)*, pages 124- 133, 2009.
- [15] G. Kurian, C. Sun, C.-H. Chen, J. Miller, J. Michel, L. Wei, D. Antoniadis, L.-S. Peh, L. Kimerling, V. Stojanovic, and A. Agarwal. Cross-layer energy and performance evaluation of a nanophotonic many-core processor system using real application workloads. In *26th IEEE International Parallel Distributed Processing Symposium*, 2012.
- [16] S. Li, J. H. Ahn, R. D. Strong, J. B. Brockman, D. M. Tullsen, and N. P. Jouppi. Mcpat: an integrated power, area, and timing modeling framework for multicore and manycore architectures. In *Proceedings of the 42nd IEEE/ACM Annual International Symposium on Microarchitecture, MICRO-42*, pages 469- 480, 2009.
- [17] S. Liu, B. Leung, A. Neckar, S. O. Memik, G. Memik, and N. Hardavellas. Hardware/software techniques for dram thermal management. In *Proceedings of the 2011 IEEE 17th International Symposium on High Performance Computer Architecture*, pages 515- 525, 2011.
- [18] S. Manipatruni, R. K. Dokania, B. Schmidt, N. Sherwood-Droz, C. B. Poitras, A. B. Apsel, and M. Lipson. Wide temperature range operation of micrometer-scale silicon electro-optic modulators. *Opt. Lett.*, 33(19):2185- 2187, Oct 2008.
- [19] B. Mondal, P. Basu, B. Reddy, H. Saha, P. Bhattacharya, and C. Roychoudhury. Oxidized macro porous silicon layer as an effective material for thermal insulation in thermal effect microsystems. In *International Conference on Emerging Trends in Electronic and Photonic Devices Systems*, pages 202- 206, Dec 2009.
- [20] C. Nitta, M. Farrens, and V. Akella. Addressing system-level trimming issues in on-chip nanophotonic networks. In *17th IEEE International Symposium on High Performance Computer Architecture*, 2011.
- [21] K. Padmaraju and K. Bergman. Resolving the thermal challenges for silicon microring resonator devices. *Nanophotonics*, 3(4-5):269- 281, September 2013.
- [22] Y. Pan, J. Kim, and G. Memik. Flexishare: Channel sharing for an energy-efficient nanophotonic crossbar. In *Proceedings of the IEEE International Symposium on High-Performance Computer Architecture*, 2010.
- [23] Y. Pan, J. Kim, and G. Memik. Featherweight: low-cost optical arbitration with qos support. In *Proceedings of the 44th IEEE/ACM Annual International Symposium on Microarchitecture*, pages 105- 116, 2011.
- [24] Y. Pan, P. Kumar, J. Kim, G. Memik, Y. Zhang, and A. Choudhary. Firefly: Illuminating future network-on-chip with nanophotonics. In *Proceedings of the 36th Annual International Symposium on Computer Architecture*, 2009.
- [25] P. Rosenfeld, E. Cooper-Balis, and B. Jacob. Dramsim2: A cycle accurate memory system simulator. *Computer Architecture Letters*, 10(1):16- 19, 2011.
- [26] M. M. Sabry, A. K. Coskun, D. Atienza, T. S. Rosing, and T. Brunschweiler. Energy-efficient multiobjective thermal control for liquid-cooled 3-d stacked architectures. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 30(12):1883- 1896, 2011.
- [27] K. Sankaranarayanan, B. H. Meyer, W. Huang, R. Ribando, H. Haj-Hariri, M. R. Stan, and K. Skadron. Architectural implications of spatial thermal filtering. *Integration VLSI Journal*, 46(1):44- 56, Jan. 2013.
- [28] T. Sarvey, Y. Zhang, Y. Zhang, H. Oh, and M. Bakir. Thermal and electrical effects of staggered micropin-fin dimensions for cooling of 3d microsystems. In *IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm)*, 2014.
- [29] K. Skadron, M. R. Stan, W. Huang, S. Velusamy, K. Sankaranarayanan, and D. Tarjan. Temperature-aware microarchitecture. In *Proceedings of the Annual International Symposium on Computer Architecture*, 2003.
- [30] C. Sun, C.-H. O. Chen, G. Kurian, L. Wei, J. Miller, A. Agarwal, L.-S. Peh, and V. Stojanovic. Dsent - a tool connecting emerging photonics with electronics for opto-electronic networks-on-chip modeling. In *6th IEEE/ACM International Symposium on Networks-on-Chip*, 2012.
- [31] D. Vantrease, N. L. Binkert, R. Schreiber, and M. H. Lipasti. Light speed arbitration and flow control for nanophotonic interconnects. In *Proceedings of the 42nd IEEE/ACM Annual International Symposium on Microarchitecture*, pages 304- 315, 2009.
- [32] D. Vantrease, R. Schreiber, M. Monchiero, M. McLaren, N. P. Jouppi, M. Fiorentino, A. Davis, N. Binkert, R. G. Beausoleil, and J. H. Ahn. Corona: System implications of emerging nanophotonic technology. In *Proceedings of the 35th Annual International Symposium on Computer Architecture*, pages 153- 164, 2008.
- [33] T. F. Wenisch, R. E. Wunderlich, M. Ferdman, A. Ailamaki, B. Falsafi, and J. C. Hoe. SimFlex: statistical sampling of computer system simulation. *IEEE Micro*, 26(4):18- 31, Jul-Aug 2006.
- [34] T. Zhang, J. Abellan, A. Joshi, and A. Coskun. Thermal management of manycore systems with silicon-photonic networks. In *Design, Automation and Test in Europe Conference and Exhibition (DATE)*, March 2014.
- [35] Y. Zhang, H. Oh, and M. Bakir. Within-tier cooling and thermal isolation technologies for heterogeneous 3d ics. In *2013 IEEE International 3D Systems Integration Conference (3DIC)*, pages 1- 6, Oct 2013.
- [36] J. Zhao, X. Dong, and Y. Xie. Cost-aware three-dimensional (3d) many-core multiprocessor design. In *47th ACM/IEEE Design Automation Conference, DAC-2010*, June 2010.
- [37] L. Zhou and A. Kodi. Probe: Prediction-based optical bandwidth scaling for energy-efficient nocs. In *Seventh IEEE/ACM International Symposium on Networks on Chip (NoCS)*, pages 1- 8, 2013.
- [38] W. Zortman, A. Lentine, D. Trotter, and M. Watts. Integrated cmos compatible low power 10gbps silicon photonic heater-modulator. In *National Fiber Optic Engineers Conference and Optical Fiber Communication Conference and Exposition (OFC/NFOEC)*, pages 1- 3, March 2012.