

WAVES: Wavelength Selection for Power-Efficient 2.5D-Integrated Photonic NoCs

Aditya Narayan*, Yvain Thonnart[†], Pascal Vivet[†], César Fuguet Tortolero[†] and Ayse K. Coskun*

*Boston University, Boston, MA 02215, USA; E-mail: {adityan, acoskun}@bu.edu

[†]Université Grenoble Alpes, CEA-Leti, Grenoble, France; E-mail: {first.last}@cea.fr

Abstract—Photonic Network-on-Chips (PNoCs) offer promising benefits over Electrical Network-on-Chips (ENoCs) in many-core systems owing to their lower latencies, higher bandwidth, and lower energy-per-bit communication with negligible data-dependent power. These benefits, however, are limited by a number of challenges. Microring resonators (MRRs) that are used for photonic communication have high sensitivity to process variations and on-chip thermal variations, giving rise to possible resonant wavelength mismatches. State-of-the-art microheaters, which are used to tune the resonant wavelength of MRRs, have poor efficiency resulting in high thermal tuning power. In addition, laser power and high static power consumption of drivers, serializers, comparators, and arbitration logic partially negate the benefits of the sub-pJ operating regime that can be obtained with PNoCs. To reduce PNoC power consumption, this paper introduces WAVES, a wavelength selection technique to identify and activate the minimum number of laser wavelengths needed, depending on an application’s bandwidth requirement. Our results on a simulated 2.5D manycore system with PNoC demonstrate an average of 23% (resp. 38%) reduction in PNoC power with only <1% (resp. <5%) loss in system performance.

I. INTRODUCTION

The core count in manycore systems is increasing to support extremely parallel abundant-data applications with higher compute performance and improved energy efficiency needs [1]. This dense integration of cores results in larger die sizes, which cause lower yield and, as a result, cost challenges. 2.5D integration of smaller chiplets on a large interposer chip has the potential to achieve a higher throughput per watt (or per volume) at a higher manufacturing yield than a large 2D chip [2], [3]. However, with a large core count and longer distances between chiplets, the system performance and energy efficiency is then bottlenecked by the Network-on-Chip (NoC) latencies and bandwidth.

Photonic NoCs (PNoCs) with wavelength-division multiplexing are currently emerging as promising alternatives over Electrical NoCs (ENoCs). In PNoCs, multiple optical signals can be multiplexed onto the same waveguide using microring resonators (MRR), thereby enabling a high internal bandwidth [4]–[6]. Numerous works demonstrating the feasibility of integrating photodiodes [7], low-loss waveguides [8], grating couplers [9], and MRR modulators and filters [10] through slightly adapted or unmodified CMOS process have paved the way for efficient realization of PNoCs.

Figure 1 shows an example PNoC. An off-chip laser source emits multiple optical signals that are coupled onto an on-chip waveguide using grating couplers. MRRs that are fabricated with particular dimensions to resonate at desired laser wavelengths perform data modulation at the *transmit* site (Tx) as well as data filtering at the *receive* site (Rx). An MRR modulates the serialized data at Tx onto a laser wavelength, by resonating at the same wavelength as the laser. The optical signal travels through the on-chip waveguide and is filtered at Rx by an MRR that is also resonating at the same wavelength. Based on the intensity of the detected optical signal by the photodetector, a transimpedance amplifier (TIA) feeds the output electrical signal to Rx comparators that deserialize the data.




Fig. 1: An example PNoC.

However, a major factor hampering the maturity of PNoCs is the significant power overhead of the lasers, drivers, serializers, comparators and arbitration logic that increases linearly with the number of laser wavelengths in the system [11], [12]. Furthermore, the MRRs undergo resonant wavelength shifts due to on-chip thermal variations (TV) and manufacturing process variations (PV). The MRRs are typically supplied with heating power to tune them back to desired laser wavelengths and compensate these TV- and PV-induced wavelength shifts. This heating power, however, can get as high as 30% of the system power [11], thereby adversely impacting the promises of sub-pJ advantages of silicon PNoC technology.

Recent design-level innovations show that an analog thermal control loop can remap and lock the MRRs to laser wavelengths compensating for TV- and PV-induced wavelength shifts within $100\mu s$ [13]. This thermal tuning interval is greater than the laser power on/off latency ($<5ns$ with $<100pm$ transient error [14]), which enables us to selectively activate laser wavelengths. In our work, we demonstrate that the bandwidth requirements for different applications vary substantially and, therefore, activating all the available laser wavelengths results in increased PNoC power with minimal performance improvement. These observations motivate the need to identify and activate a reduced number of laser wavelengths based on an application’s bandwidth requirements. Since the MRR lock latency dominates laser switching time, our technique provides a low-latency wavelength selection to reduce the PNoC power. Our major contributions are as follows:

- 1) We design a wavelength selection technique, WAVES, by accounting for the on-chip TV and PV. Based on the bandwidth requirement of an application, WAVES identifies the minimum number of laser wavelengths and activates the best combination resulting in reduced PNoC power consumption with minimal losses in performance.
- 2) We develop a cross-layer simulation framework to model the system performance and PNoC power (laser, electronics and heating) under different activated laser wavelengths and MRR lock status, which is a strong function of on-chip TV and PV. Through this framework, we explore the optimizations of PNoC power arising from the device-level MRR locking under different system-level constraints (DVFS, thread combinations), based on technology parameters measured on silicon. Our simulation results on a 96-core 2.5D system with PNoC demonstrate 23% (resp. 38%, 42%) reduction in PNoC power with only 1% (resp. 5%, 10%) performance loss.

II. RELATED WORK

Several prior approaches address the challenge of designing energy-efficient PNoCs. Bahadori *et al.* explore the tradeoff between an increased number of channels at a high bitrate for maximum aggregated bandwidth and the energy consumption of a PNoC [15]. Luo *et al.* propose an optimized wavelength allocation by mapping an application task graph onto the system architecture, thereby reducing laser power arising from increased bit error rate [16]. In contrast, our proposed wavelength selection addresses the tradeoff between system performance and PNoC power by activating a sufficient number of laser wavelengths, depending on an application’s bandwidth needs.

A primary component of PNoC power is the high heating power overhead required to thermally tune the MRRs to the laser wavelength and to overcome the TV and PV. Major efforts to improve the efficiency of thermal heaters at the device-level include full-FSR tuning range [17], polymer material coating to prototype athermal materials [18], and substrate removal underneath the MRRs for ultralow tuning power [19].

However, most of the efforts focusing on PNoC design is incognizant to the running application or the system architecture, indicating an opportunity for optimization at a system-level. There have been prior efforts to manage thermal gradients around communicating MRRs using temperature-aware thread allocation and migration. *RingAware* [20] and *FreqAlign* [21] are workload allocation policies that balance temperature around MRRs to reduce the thermal tuning range. *Aurora* [22] encompasses a cross-layer approach to reduce thermal tuning power by applying bias current to MRRs to control small TVs, performing DVFS and packet routing around hotter regions and running a scheduling policy for the outer cores.

A common drawback of these techniques is the performance overhead due to thread migration, scheduling or packet-routing. Our wavelength selection technique incurs low latency with minimal storage overhead. To our knowledge, our work is the first to provide a detailed account of PNoC power for a 2.5D manycore system with PNoCs. Our wavelength selection approach, WAVES, is orthogonal to most prior system-level studies and can be combined with them to develop an energy-efficient PNoC.

III. 2.5D MANYCORE SYSTEMS WITH PNO C

In this section, we first present our target 2.5D manycore system with PNoC called *Processors On Photonic Silicon Interposer ARchitecture (POPSTAR)* and then provide details on computing the total PNoC power consumption (including laser, drivers, serializers, comparators, TIA and thermal tuning).

A. Architecture overview

POPSTAR is a 2.5D manycore system with a PNoC, consisting of six compute chiplets and eight TxRx chiplets stacked on a photonic interposer as shown in Fig. 2(a). The compute chiplets and the TxRx chiplets are designed using an STMICRO C28FDSOI technology, and the photonic interposer uses an optical FEOL technology. An off-chip laser emits up to six wavelengths that are carried by a vertical fiber attachment and coupled onto the waveguides in the interposer via grating couplers.

1) *Compute Chiplets*: We assume that the 96 cores in POPSTAR are organized in six 3D-ready *TSARLET* [3] compute chiplets, with each chiplet consisting of 16 cores, as shown in Fig. 2(b). In our study, each core’s architecture is similar to the IA-32 core from Intel SCC [23]. The 16 cores in a compute chiplet are organized in 4 clusters, with 4 cores per

Table I: Power consumption of different elements [11]

Component	Active Power		Idle Power	
	Notation	Value (mW)	Notation	Value (mW)
Laser	P_L	30		0
Serializer	$P_{srl,a}$	3	$P_{srl,i}$	1
Driver	P_{drv}	3		0
Rx Comparator	$P_{cmp,a}$	1	$P_{cmp,i}$	0.33
TIA	P_{TIA}	2		0
Arbitration and Flow Control	$P_{arb,a}$	32	$P_{arb,i}$	10

Table II: Notations used

Notation	Description	Value
C	Number of TxRx chiplets (and waveguides)	8
λ_{tot}	# of available laser wavelengths	6
λ_{act}	# of activated laser wavelengths	1,2, .. 6
λ_{min}	# of laser wavelengths required for an application	1,2, .. 6

cluster. Each core has a 16KB private L1 I/D cache. There is a 256KB shared distributed L2 cache per cluster, and an adaptive distributed L3 cache with four 1MB L3 tiles per chiplet.

2) *PNoC Architecture*: The on-chip PNoC handles the communication and coherence traffic between L1-L2, L2-L3, and L3-DRAM. The global network topology connecting the compute chiplets is a PNoC-based Single-Writer Multiple-Reader (SWMR) link as illustrated in Fig. 2(c). The SWMR topology is mapped onto a U-shaped spiral of waveguides on the interposer. The TxRx chiplets, described in Sec. III-A4, are responsible for the Electrical-to-Optical (E-O) and Optical-to-Electrical (O-E) conversion. Each of the six compute chiplets accesses the PNoC through a TxRx chiplet with a 96-bit wide interface operating at 750MHz. Two TxRx chiplets are used for servicing external IO and DRAM requests. Each wavelength supports a data rate of 12Gbps, resulting in a peak aggregate bandwidth of 576Gbit/s.

3) *Microring Resonator Group (MRRG)*: Each TxRx chiplet connects to an MRRG on the interposer. For every MRRG, there is one Tx waveguide and seven Rx waveguides coming from the other MRRGs forming a spiral of SWMR links as depicted in Fig. 2(c). All six wavelengths are evenly spaced in a free spectral range (FSR) of $10.8nm$, resulting in a $1.8nm$ spacing between adjacent laser wavelengths, around a center wavelength of $1310nm$. These wavelengths are modulated by Tx MRRs, and the destination TxRx chiplet filters those signals using six of its 42 Rx MRRs. We assume that all the 48 MRRs in an MRRG are normally maintained slightly off-resonance by the TxRx chiplet so that they are immediately responsive when required for communication [13].

4) *TxRx Chiplets*: A TxRx chiplet consists of the electronic circuitry for E-O and O-E conversion, as shown in Fig. 2(d). For every Tx MRR in an MRRG, there is a 16:1 serializer and a modulation driver. Similarly, for every Rx MRR in an MRRG, there is a filter bias, a TIA and a comparator. We assume that an analog thermal control loop (as shown in recent work [13]) detects the photodiode current and applies the heating power to thermally tune and lock the MRRs to the nearest laser wavelength. Finally, FIFO queues and multiplexers handle the flow control and communicate with the compute chiplets using local 2.5D passive connections.

B. Power consumption of the PNoC

The power consumption of a PNoC includes laser, serializers, drivers, comparators, TIAs, arbitration and flow control, and thermal tuning and control. Table I displays the active and idle power of different elements in a PNoC based on post-layout simulations of the TxRx chiplet. Table II shows the notations used in our work.

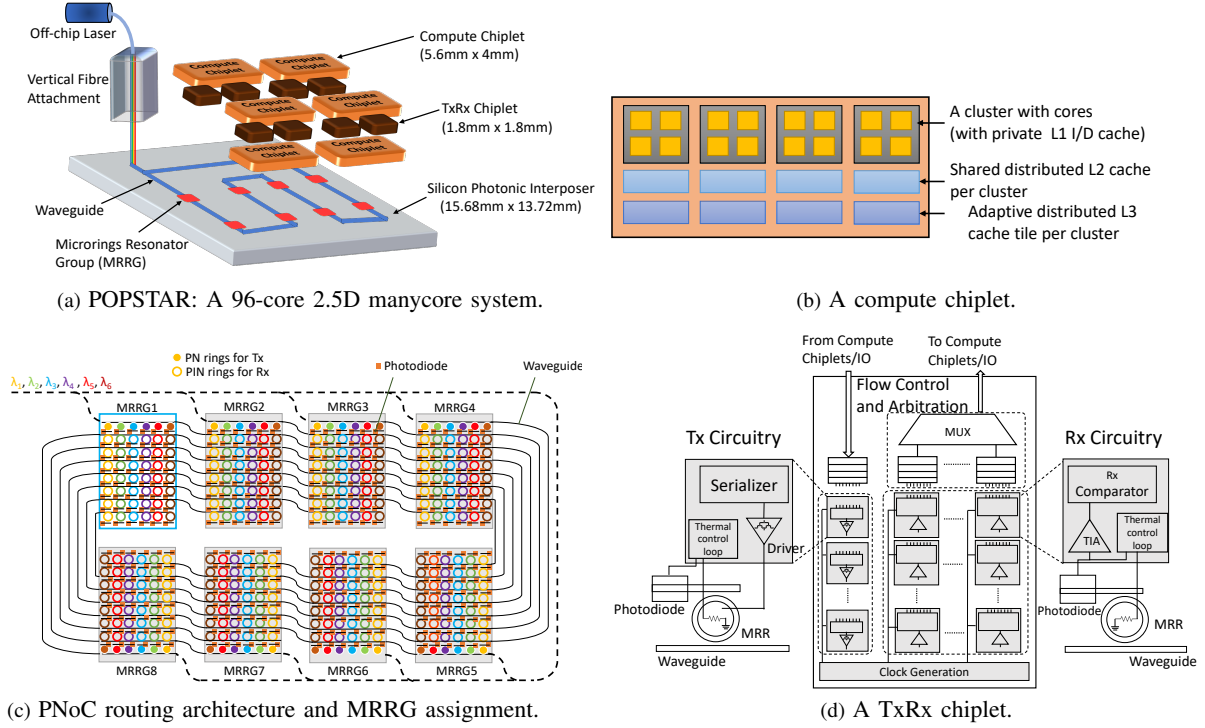


Fig. 2: The POPSTAR architecture.

The laser power is determined by taking into account the wall-plug efficiency and all the losses in the waveguides, couplers, filters and modulators. The overall **laser power**, P_{laser} , can be expressed as:

$$P_{laser} = P_L \cdot C \cdot \lambda_{act} \quad (1)$$

The first stage in a PNoC is serialization, where the data is serialized to be modulated onto the MRR. To obtain an acceptable extinction ratio at the MRR, double-rail drivers with high voltage-swings ($> V_{DD}$) are typically used [11]. For transmitting data over the six available laser wavelengths, all the six serializers and six drivers within all the eight TxRx chiplet are active. On the Rx side, active channels are biased to filter the wavelengths. TIA amplifies the detected electrical signal from the photodiode and feeds the value to the Rx comparators, where the data is deserialized. FIFO queues and a multiplexer handle the arbitration and flow control. The serializers, comparators and arbiters are clocked for precise timing control and consume idle power. Depending on λ_{act} , the overall **electronics power** of all TxRx chiplets, P_{EOE} , is determined as follows:

$$P_{Tx} = P_{drv} \cdot \lambda_{act} + P_{srl,a} \cdot \lambda_{act} + P_{srl,i} \cdot (\lambda_{tot} - \lambda_{act}) \quad (2)$$

$$P_{Rx} = P_{TIA} \cdot \lambda_{act} + P_{cmp,a} \cdot \lambda_{act} + P_{cmp,i} \cdot (\lambda_{tot} \cdot C - \lambda_{act}) \quad (3)$$

$$P_{arb} = P_{arb,a} \cdot \frac{\lambda_{act}}{\lambda_{tot}} + P_{arb,i} \cdot \frac{\lambda_{tot} - \lambda_{act}}{\lambda_{tot}} \quad (4)$$

$$P_{EOE} = C \cdot (P_{Tx} + P_{Rx} + P_{arb}) \quad (5)$$

Another major contributor to PNoC power is the thermal tuning power. Since the resonant wavelength of MRRs is highly susceptible to TV and PV, the MRRs experience a wavelength shift from the designed laser wavelength. Therefore, the MRRs need to be thermally tuned to the nearest laser wavelength. This is conventionally achieved by heat injection via Joule effect using resistive heaters inside the MRRs,

thereby increasing the MRR wavelength to the nearest laser wavelength. In our system, we assume that the MRRs have a thermal sensitivity ($\frac{d\lambda}{dT}$) of 78pm/K and the heaters have an efficiency ($\frac{d\lambda}{dH}$) of 120pm/mW [13]. Given the small area footprint of an MRRG, we assume that the steady state temperature over a single MRRG is uniform, so all the MRRs in an MRRG undergo the same TV-induced wavelength shift. However, all the eight MRRGs are possibly at different temperatures, owing to variable chip activity. The overall wavelength shift ($\Delta\lambda_{shift}$) for an MRR is given by Eq. (6), where ΔT is the difference between the MRRG temperature and the ambient temperature and $\Delta\lambda_{shift,PV}$ is the PV-induced wavelength shift:

$$\Delta\lambda_{shift} = \frac{d\lambda}{dT} \cdot \Delta T + \Delta\lambda_{shift,PV} \quad (6)$$


In a single MRRG, one Tx MRR and seven Rx MRR are heated for every activated laser wavelength. The total **heating power**, P_{heat} , is calculated by aggregating the heating power of the MRRs over all the MRRGs and is given by Eq. (8), where $\Delta\lambda_{heat}$ is the required wavelength shift to the nearest laser wavelength for an MRR:

$$\Delta\lambda_{heat} = \frac{FSR}{\lambda_{tot}} - (\Delta\lambda_{shift} \bmod \frac{FSR}{\lambda_{tot}}) \quad (7)$$

$$P_{heat} = \sum_{i=1}^C \sum_{r=1}^C \cdot C \cdot \lambda_{act} \cdot \frac{\Delta\lambda_{heat,ir}}{\frac{d\lambda}{dH}} \quad (8)$$

IV. WAVES: WAVELENGTH SELECTION FOR PNOCS

As evident from the previous section, the power consumption (P_{laser} , P_{EOE} and P_{heat}) along a PNoC increases with λ_{act} . In this section, we first motivate the need for wavelength selection and then describe how the MRRs lock to the nearest laser wavelength with TV- and PV-induced wavelength shifts.



(a) Normalized execution time. (b) System power.
Fig. 3: Performance and power breakdown with different λ_{act} .

A. Applications' bandwidth needs

Figure 3(a) illustrates the system performance of selected applications from SPLASH2 [24] and PARSEC [25] benchmark suites at different inter-chiplet bandwidth settings, i.e., λ_{act} . Our simulation framework to run experiments is presented in Sec. V. We plot the normalized execution time (normalized to $\lambda_{act} = 6$) for increasing λ_{act} . We observe a general trend that the system performance tends to saturate at a particular λ_{act} . For certain applications such as *cholesky* and *swaptions*, a larger λ_{act} is desirable for high performance. On the other hand, for applications such as *lu.cont* and *barnes*, a lower λ_{act} provides similar performance as λ_{tot} . Figure 3(b) shows the power breakdown of using different λ_{act} . It is evident that the power consumption increases with increasing λ_{act} . Since the performance saturates at a particular λ_{act} , which is dependent on the application's bandwidth requirement, activating all the available λ_{tot} laser wavelengths may increase the power without providing notable performance improvements.

To exploit this observation, we perform a design space exploration (DSE) to determine the minimum required number of laser wavelengths (λ_{min}) that is sufficient to cater to the bandwidth requirement of an application. We set a performance loss threshold (L_{thr}) that is deemed acceptable for a system. For each application, we identify the λ_{act} where the system performance loss is under L_{thr} as compared to the performance with λ_{tot} . We define this value of λ_{act} as λ_{min} . Once we determine the λ_{min} , we then need to activate the best combination of λ_{min} laser wavelengths that consume the lowest PNoC power by accounting for TV and PV. Section IV-B accounts for only TV-induced wavelength shifts and explains the MRR locking mechanism. Section IV-C incorporates PV in addition to TV and activates the best λ_{min} combination that results in the lowest thermal tuning power.

B. Accounting for TV-induced wavelength shifts

Figure 4 illustrates the designed and shifted resonant wavelength due to TV and PV for six Tx MRRs of an MRRG. In an ideal case with no PV, the six MRRs resonate at the six laser wavelengths as seen in Fig. 4(a). From our DSE, we determine the number of laser wavelengths for an application to be λ_{min} (e.g., $\lambda_{min}=2$ in Fig. 4). Due to chip activity and the resulting TV, the resonant wavelength of MRRs shift as illustrated in Fig. 4(b). Since the MRRs are evenly spaced within the FSR, this TV-induced resonance shift can be overcome by applying the same amount of heating at each MRR and, therefore, any λ_{min} laser wavelengths can be activated. For example, we can activate the first λ_{min} wavelengths in order to compensate this TV-induced wavelength shift. The thermal control loop in the TxRx chiplet supplies heating power to λ_{tot} MRRs until λ_{min} MRRs lock to the activated laser wavelengths. The other MRRs, even though heated at maximum, cannot lock, and the thermal control loop removes their heating power, as shown in Fig. 4(b). The locked MRRs for Tx and Rx are used for

communication at the activated laser wavelengths. As MRRGs can be at different temperatures due to uneven chip activity, the set of MRRs locked to the laser wavelengths can be different among MRRGs. Thus, the λ_{min} selection process is performed dynamically based on MRR tuning ranges and lock status.

C. Accounting for TV- and PV-induced wavelength shifts

Due to PV, the MRRs do not resonate at the design wavelength and encounter an additional wavelength shift. We model the local MRRG PV as a gaussian distribution with a standard deviation of 100pm .

Figure 4(c) illustrates the combined TV- and PV-induced resonant wavelength shifts of MRRs. These shifts are different for all the MRRs within an MRRG, and hence not evenly spaced within the FSR. Therefore, all the MRRs in an MRRG now require variable amount of heating. In this case, activating the first 2 laser wavelengths as claimed in Sec. IV-B results in a larger tuning range and thereby, a suboptimal activation of laser wavelengths as depicted in Fig. 4(c). Figure 4(d) demonstrates that activating laser wavelengths λ_3 and λ_4 results in a lower tuning range. Hence, the combined TV and PV effect forces the need to activate the best combination of λ_{min} for a reduced tuning range and lower heating power.

D. System integration of WAVES

To implement WAVES on a simulated real system prototype and select the best combination of λ_{min} , we implement the following algorithm: (1) store the PV-induced resonant wavelength shift of all 48 MRRs in all 8 MRRGs as an $8 \times 8 \times 6$ table; (2) determine all possible combinations of activating λ_{act} as $\binom{\lambda_{tot}}{\lambda_{act}}$ (≤ 20 combinations for $\lambda_{tot} = 6$); (3) for each combination of λ_{act} , store the heating power required in all the MRRGs for the operating temperature range (e.g. 300K-380K) in a lookup table (LUT); (4) given the λ_{min} determined from our DSE for an application profile, and a temperature estimate or measurement at each MRRG, compute the heating power for all combinations, $\binom{\lambda_{tot}}{\lambda_{min}}$, by accessing the LUT; (5) determine the minimum heating power and activate the corresponding λ_{min} laser wavelengths for an application profile. Steps (1)-(3) are performed at design time, and steps (4) and (5) are performed at runtime.

We compute the memory requirement of the LUT as 200KB so it can be easily stored on-chip. The MRR lock-time to the laser wavelength is $100\mu\text{s}$ [13]. We determine the LUT access time as $\leq \binom{\lambda_{tot}}{\lambda_{tot}/2} \cdot C$ lookups and additions (160 cycles for $C = 8$ and $\lambda_{tot} = 6$), while laser power-on latency from a cold state is 5ns [14], hence each takes much less than $100\mu\text{s}$. Thus, the latency overhead of dynamic laser activation is negligible compared to MRR lock time, and can be run periodically or every time the MRRs get unlocked due to high temperature drift. As WAVES does not depend on specific architecture parameters, our technique is scalable and can be extended to larger systems with a higher number of MRRs.

V. SIMULATION FRAMEWORK

To evaluate the impact of WAVES, we set up a simulation framework consisting of a performance simulator, core power and PNoC power model, and a thermal simulator as shown in Fig. 5. For performance simulation, we model the POPSTAR architecture in Sniper [26]. We use multithreaded applications from SPLASH-2 [24] and PARSEC [25] benchmark suites. We experiment with two different voltage/frequency (V/f) settings and four thread combinations for each of the applications. We activate different λ_{act} to simulate varying internal bandwidth. Table III details our experiments. For each experiment, we execute 10 billion instructions in the parallel region or the full region of interest if it finishes earlier.

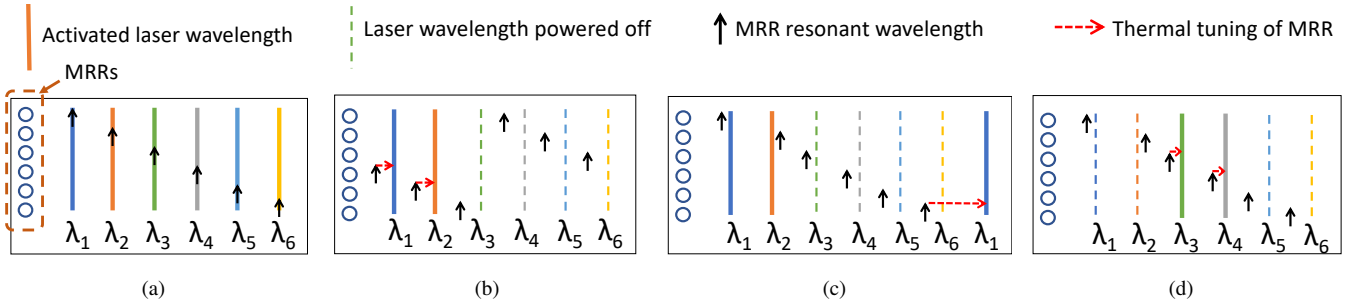


Fig. 4: Resonant wavelengths for a given MRRG: (a) Design intent, (b) Ideal locking on first $\lambda_{min} = 2$ wavelengths after TV-induced shift, (c) Suboptimal locking by activating first $\lambda_{min}=2$ with TV and PV, (d) Locking by activating best $\lambda_{min}=2$ with TV and PV.





Fig. 5: Simulation framework.

Table III: List of experiments

Experiment	Detail
Applications	
<i>high_comm</i>	canneal, swaptions, cholesky
<i>low_comm</i>	blackscholes, barnes, lu.cont
DVFS setting	
<i>high_perf</i>	V=0.95V, f=800MHz; avg core power=1.05W
<i>low_perf</i>	V=0.85V, f=533MHz; avg core power=0.83W
Thread combinations	
$24 \cdot i$ threads	i active clusters in each chiplet ($i=1,2,3,4$)

We use McPAT [27] to compute the core and cache power consumption. We collect power traces from McPAT simulations and scale the raw power data to the published power values for the Intel SCC cores. We use Equations 1-5 to compute the laser and electronics power depending on λ_{act} .

We use the 3D extension [28] of HotSpot [29] to determine the steady-state temperature of all the MRRGs. In order to obtain an accurate thermal map, we calibrate HotSpot temperatures to temperatures obtained from Project Sahara [30], which is a sign-off thermal tool from Mentor, a Siemens Business, for simulating detailed 3D circuits within its package and board. We obtain HotSpot temperatures within 2% error margin of Project Sahara on average. Figure 6 illustrates the thermal map in Project Sahara and HotSpot. We assume that the MRRs are designed to resonate at the laser wavelengths at an ambient temperature of 300K.



(a) Interposer layer.

(b) Chiplet layer.

Fig. 6: Thermal map in Sahara and HotSpot.

VI. EXPERIMENTAL RESULTS

To investigate the power savings obtained from WAVES, we conduct experiments on applications with different communication intensities under varying system settings (DVFS

and thread combinations) and demonstrate their combined impact on λ_{min} . Our baseline system activates all the available laser wavelengths ($\lambda_{act} = \lambda_{tot}$). We experiment with three different L_{thr} values (1%, 5% and 10%) to determine the λ_{min} ($\lambda_{act} = \lambda_{min}$). For each L_{thr} , we compute the PNoC power by activating the first λ_{min} laser wavelengths (Sec. IV-B), and by activating the best combination of λ_{min} laser wavelengths (Sec. IV-C). Figure 7 shows the PNoC power savings obtained with WAVES for different applications under varying system settings. We tabulate our average power savings in Table IV.

For *low_comm* applications, the system performance saturates for a low λ_{act} , as seen in Fig. 3(a). As a consequence, even for a 1% L_{thr} , a lower λ_{min} is activated and we observe average power savings of 38%. However, for *high_comm* applications, the high network traffic demands higher λ_{min} , resulting in average power savings of 8% for a 1% L_{thr} .

Subsequently, we experiment with two DVFS settings to evaluate the effect of compute performance on WAVES. We observe that *high_perf* (Figs. 7(e)-7(h)) desires a higher λ_{min} for the same L_{thr} than *low_perf* (Figs. 7(a)-7(d)), and therefore, result in lower power savings. On average, we obtain 19% and 26% power savings for *high_perf* and *low_perf* respectively. The lower power savings in *high_perf* is attributed to the larger on-chip thermal gradient due to increased logic power, giving rise to higher thermal tuning power.

Next, we run applications with different threads counts and combinations to study their effect on WAVES. Figures 7(a)-7(d) (and 7(e)-7(h)) show the power savings for different applications running with 24, 48, 72 and 96 threads respectively. In most cases, larger thread counts result in increased inter-chiplet network traffic among the communicating threads and result in a higher λ_{min} . This is evident in *high_comm* applications running 96 threads, particularly *canneal* and *cholesky* in Fig. 7(d) and Fig. 7(h), which require all six laser wavelengths to be activated even for an L_{thr} of 10%.

WAVES achieves 23% (resp. 38%, 42%) average PNoC power savings with only 1% (resp. 5%, 10%) performance loss. In addition, we demonstrate that activating the best combination of these laser wavelengths is feasible with negligible storage and latency overhead.

Table IV: Average power savings over different configurations

Configuration	Setting	1%	L_{thr} 5%	10%
Applications	<i>high_comm</i>	8%	21%	26%
	<i>low_comm</i>	38%	56%	61%
DVFS settings	<i>high_perf</i>	19%	34%	39%
	<i>low_perf</i>	26%	41%	45%
Thread combinations	24 threads	28%	48%	54%
	48 threads	18%	37%	45%
	72 threads	26%	34%	36%
	96 threads	18%	31%	33%
Overall		23%	38%	42%

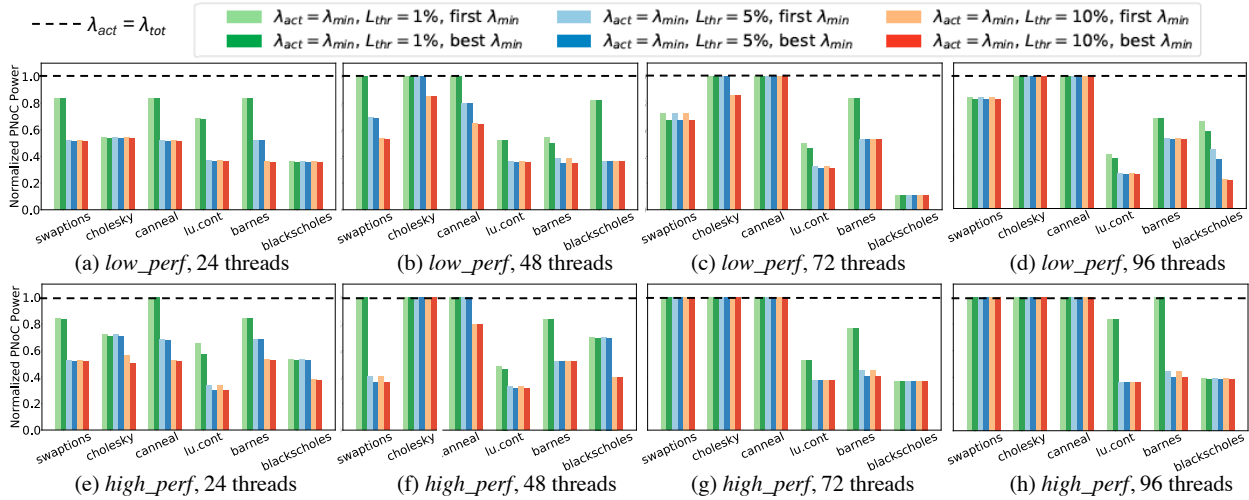


Fig. 7: PNoC power savings when using WAVES under different DVFS settings and thread combinations. The six bars for each application shows power savings obtained by activating first λ_{min} and the best combination of λ_{min} laser wavelengths for three L_{thr} options. The baseline case (horizontal line) activates all λ_{tot} laser wavelengths.

VII. CONCLUSIONS AND FUTURE WORK

PNoCs are promising alternatives for ENoCs, however, the practical integration of PNoCs in manycore systems is contingent on improving their energy efficiency. MRRs are highly susceptible to on-chip TV and PV, resulting in high PNoC power. Our wavelength selection technique, WAVES, accounts for these variations when selecting a minimum number of laser wavelengths and provides 23% average power savings under 1% loss in system performance. We demonstrate the feasibility of WAVES on a simulated 2.5D-integrated PNoC manycore system with a detailed account of MRR wavelength locking under TV and PV. Our work is orthogonal to most other design or runtime optimization methods for PNoCs and can be applied in tandem. WAVES can be further improved by using an online mechanism based on network and memory accesses to predict the bandwidth requirement for an application.

ACKNOWLEDGEMENT

We thank Lee Wang for her help in Project Sahara thermal simulations. This work was funded partly by the CARNOT institute and the French National Program Programme d'Investissements d'Avenir, IRT Nanoelec under Grant ANR-10-AIRT-05. This work was also funded partly by NSF CCF-1716352.

REFERENCES

- [1] M. Kistler, M. Perrone, and F. Petrini, "Cell multiprocessor communication network: Built for speed," *IEEE micro*, no. 3, 2006.
- [2] F. Eris *et al.*, "Leveraging thermally-aware chiplet organization in 2.5D systems to reclaim dark silicon," in *Proc. Design, Automation and Test in Europe (DATE)*, 2018.
- [3] E. Guthmuller *et al.*, "A 29 Gops/Watt 3D-ready 16-core computing fabric with scalable cache coherent architecture using distributed L2 and adaptive L3 caches," in *Proc. European Solid-State Circuits Conf.*, 2018.
- [4] A. Shacham, K. Bergman, and L. P. Carloni, "Photonic networks-on-chip for future generations of chip multiprocessors," *IEEE Trans. on Computers*, vol. 57, no. 9, 2008.
- [5] A. Joshi *et al.*, "Silicon-photonics networks for global on-chip communication," in *Proc. Int. Symp. on Networks-on-Chip (NoCS)*, 2009.
- [6] D. Vantrease *et al.*, "Corona: System implications of emerging nanophotonic technology," in *Proc. Int. Symp. on Computer Architecture (ISCA)*, 2008.
- [7] L. Virost *et al.*, "Germanium avalanche receiver for low power interconnects," *Nature communications*, vol. 5, 2014.
- [8] J. Cardenas *et al.*, "Low loss etchless silicon photonic waveguides," *Optics express*, vol. 17, no. 6, 2009.
- [9] M. T. Wade *et al.*, "75% efficient wide bandwidth grating couplers in a 45 nm microelectronics CMOS process," in *Proc. IEEE Optical Interconnects Conference (OI)*, 2015.
- [10] W. Bogaerts *et al.*, "Silicon microring resonators," *Laser & Photonics Reviews*, vol. 6, no. 1, 2012.

- [11] R. Polster *et al.*, "Efficiency optimization of silicon photonic links in 65-nm CMOS and 28-nm FDSOI technology nodes," *IEEE Trans. on Very Large Scale Integration Systems*, vol. 24, no. 12, 2016.
- [12] M. Bahadori *et al.*, "Energy-bandwidth design exploration of silicon photonic interconnects in 65nm CMOS," in *Proc. OI*, 2016.
- [13] Y. Thonnart *et al.*, "A 10Gb/s Si-photonics transceiver with 150 μ W 120 μ s-lock-time digitally supervised analog microring wavelength stabilization for 1Tb/s/mm² die-to-die optical networks," in *Proc. Int. Solid-State Circuits Conference (ISSCC)*, 2018.
- [14] G. Simon *et al.*, "Experimental demonstration of low cost wavelength drift mitigation for TWDM systems," in *Proc. European Conference on Optical Communication (ECOC)*, 2016.
- [15] M. Bahadori *et al.*, "Energy-performance optimized design of silicon photonic interconnection networks for high-performance computing," in *Proc. DATE*, 2017.
- [16] J. Luo *et al.*, "Offline optimization of wavelength allocation and laser power in nanophotonic interconnects," *ACM Journal on Emerging Technologies in Computing Systems (JETC)*, vol. 14, no. 2, 2018.
- [17] F. Gan *et al.*, "Maximizing the thermo-optic tuning range of silicon photonic structures," in *Proc. Photonics in Switching*, 2007.
- [18] S. S. Djordjevic *et al.*, "CMOS-compatible, athermal silicon ring modulators clad with titanium dioxide," *Optics Express*, vol. 21, no. 12, 2013.
- [19] P. Dong *et al.*, "Thermally tunable silicon racetrack resonators with ultralow tuning power," *Optics express*, vol. 18, no. 19, 2010.
- [20] T. Zhang, J. L. Abellán, A. Joshi, and A. K. Coskun, "Thermal management of manycore systems with silicon-photonics networks," in *Proc. DATE*, 2014.
- [21] J. L. Abellán *et al.*, "Adaptive tuning of photonic devices in a photonic NoC through dynamic workload allocation," *IEEE Trans. on Computer-Aided Design of Integrated Circuits and Systems*, vol. 36, no. 5, 2017.
- [22] Z. Li *et al.*, "Aurora: A cross-layer solution for thermally resilient photonic network-on-chip," *IEEE Trans. on Very Large Scale Integration Systems*, vol. 23, no. 1, 2015.
- [23] J. Howard *et al.*, "A 48-core IA-32 message-passing processor with DVFS in 45nm CMOS," in *Proc. ISSCC*, 2010.
- [24] S. C. Woo *et al.*, "The SPLASH-2 programs: Characterization and methodological considerations," in *ACM SIGARCH computer architecture news*, vol. 23, no. 2, 1995.
- [25] C. Bienia, S. Kumar, J. P. Singh, and K. Li, "The PARSEC benchmark suite: Characterization and architectural implications," in *Proc. Parallel architectures and compilation techniques*, 2008.
- [26] T. E. Carlson, W. Heirman, and L. Eeckhout, "Sniper: Exploring the level of abstraction for scalable and accurate parallel multi-core simulation," in *Proc. Int. Conference for High Performance Computing, Networking, Storage and Analysis*, 2011.
- [27] S. Li *et al.*, "McPAT: an integrated power, area, and timing modeling framework for multicore and manycore architectures," in *Proc. Int. Symp. on Microarchitecture*, 2009.
- [28] J. Meng, K. Kawakami, and A. K. Coskun, "Optimizing energy efficiency of 3D multicore systems with stacked dram under power and thermal constraints," in *Proc. Design Automation Conference*, 2012.
- [29] K. Skadron *et al.*, "Temperature-aware microarchitecture," in *Proc. ISCA*, 2003.
- [30] J. Parry and L. Wang. A complete guide to 3D chip-package thermal co-design 10 key considerations. [Online]. Available: www.mentor.com