Temperature and Process Variation-Aware Wavelength Selection in 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) are promising for large-scale manycore systems due to their high bandwidth communication at sub-pJ energy-per-bit capabilities. However, such advantages come with an increased power cost, primarily from laser, electrical-optical conversion and thermal tuning of the microring resonators (MRRs). The resonant wavelength of MRRs is highly sensitive to thermal (TV) and process variations (PV), which further add to thermal tuning power. We introduce WAVES, a wavelength selection technique in PNoCs, that determines the minimum number of laser wavelengths (λ_{min}) depending on the bandwidth requirements of an application. WAVES accounts for the TV and PV of MRRs and activates the best λ_{min} that result in lowest thermal tuning power at runtime. Our simulation results on a 2.5D manycore system with PNoC demonstrate an average of 23% reduction in PNoC power with only <1% loss in system performance.**

I. Introduction

Emerging compute-heavy and data-intensive applications demand higher parallelism and larger data transfers compared to existing applications. Therefore, the performance and energy efficiency of such applications running in large manycore systems are severely hampered by network-on-chip (NoC) latencies and bandwidth. Compared to traditional electrical NoCs (ENoCs), photonic NoCs (PNoCs) have been demonstrated to provide high bandwidth at sub-pJ energy-per-bit communications [1], [2]. The growing interest towards PNoCs is further bolstered by the technological feasibility of integrating photodiodes, waveguides, couplers and MRR modulators and filters through a slightly adapted CMOS process [3], [4].

A major factor hampering the sub-pJ promises of PNoC technology is the high power resulting from lasers, electrical-optical (EO) and optical-electrical (OE) conversion. In addition, the on-chip thermal variations (TV) and manufacturing process variations (PV) induce MRR resonant wavelength shifts. The MRRs are typically supplied with heating power to tune them back to the intended laser wavelengths. The overall PNoC power also increases with the number of activated laser wavelengths (λ_{act}). A higher λ_{act} provides high bandwidth and is desirable for achieving higher performance, but the power contribution gets considerably higher.

To address such bandwidth-power tradeoffs in manycore systems with PNoCs, we argue that it is essential to determine a balanced number of laser wavelengths required for an application. Our wavelength selection technique, WAVES, determines the minimum number of laser wavelengths (λ_{min}) required for an application, given a performance loss threshold (L_{thr}). WAVES incorporates the TV and PV and activates the best λ_{min} for a given application. As the laser activation latency ($5ns$) is considerably lower than the thermal tuning control latency ($100\mu s$) [5], WAVES enables a low-overhead selective activation of laser wavelengths to reduce the PNoC power consumption.

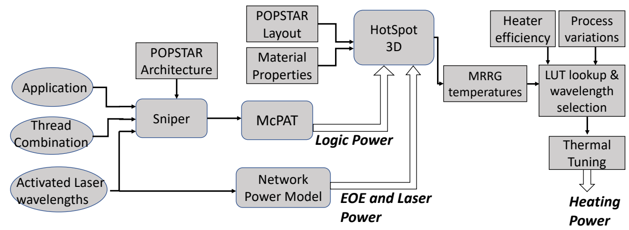

Fig. 1: Simulation framework of WAVES.

II. WAVES: Wavelength Selection in PNoCs

To demonstrate the benefits of WAVES, we conduct experiments on a 2.5D manycore system with PNoC called *Processors On Photonic Silicon inTerposer ARchitecture (POPSTAR)*. POPSTAR is a 2.5D 96-core system organized into six compute chiplets and eight TxRx chiplets. Each compute chiplet consists of 16 cores, with the core architecture similar to the IA-32 core from Intel SCC [6]. The TxRx chiplets are composed of the electronic circuitry for EO and OE conversion. The global PNoC topology is based on Single-Writer Multiple Reader links, mapped onto a U-shaped spiral of waveguides on the interposer. An off-chip laser emits photonic signals onto the interposer waveguide. The MRRs modulate and filter these photonic signals. A detailed description of *POPSTAR* is provided in recent work [7].

Figure 1 presents the simulation framework of our WAVES technique. We use Sniper [8] to simulate applications with various numbers of activated laser wavelengths. Given a performance loss threshold (L_{thr}) that is deemed acceptable for an application A, we determine the minimum number of laser wavelengths (λ_{min}) that provides system performance within L_{thr}. Activating λ_{min} at runtime is sufficient to satisfy the bandwidth requirements of the application A.

Since the on-chip TV and PV induce major wavelength shifts in the resonant wavelength of the MRRs, it is essential to incorporate the combined effects of TV and PV on the thermal tuning power. Figure 2 illustrates the designed and shifted MRR resonant wavelengths due to TV and PV. Figure 2(a) shows the design intent of the MRRs, with their resonant wavelength aligned with the laser wavelengths. Due to TV, the resonant wavelength of the MRRs within a group shift due to their thermal dependence as shown in Fig. 2(b) [1]. For an application with $\lambda_{min} = 2$, activating any two laser wavelengths results in similar thermal tuning power. However, the PV-induced wavelength shift of the MRRs is variable even with the same group, and therefore, the combined wavelength shift due to PV and TV is different for the MRRs. In this case, activating different options of λ_{min} laser wavelengths will result in different thermal tuning power. For example, activating the first 2 wavelengths as in Fig. 2(c) gives rise to

[1] We assume that the temperature of all the MRRs within a group remains same, so the TV-induced wavelength shift is the same.

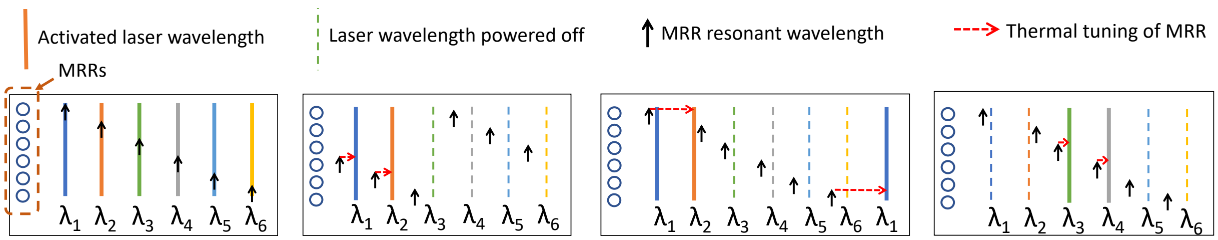

(a) (b) (c) (d)

Fig. 2: 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 λ_{min}=2 with TV and PV, (d) Locking by activating best λ_{min}=2 with TV and PV.

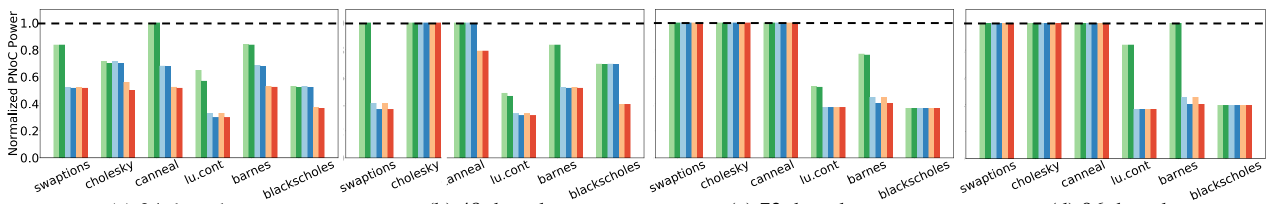

(a) 24 threads (b) 48 threads (c) 72 threads (d) 96 threads

Fig. 3: PNoC power savings when using WAVES for different 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.

a higher thermal tuning power and thereby, is a sub-optimal selection. Figure 2(d) shows that activating laser wavelengths λ_3 and λ_4 results in lowest thermal tuning range.

We use McPAT [9] to determine the core and cache power, and use an analytical model [7] to calculate the EOE and laser power for λ_{act}. Using the compute and TxRx chiplet power as inputs, we use HotSpot [10] to determine the steady-state temperatures of the MRRs. We store the PV-induced resonant wavelength shifts of all MRRs in an on-chip lookup table (LUT). At runtime, given a λ_{min} for an application, we poll the LUT to calculate the thermal tuning power of all different options for selecting λ_{min}. The λ_{min} wavelengths that corresponds to the lowest thermal tuning power is activated for that application. The memory requirement of LUT is calculated as 200KB, and the latency of dynamic laser activation is negligible compared to MRR tuning. Thus, WAVES provides a low-latency, low-storage wavelength selection for PNoCs.

III. Experimental Results

To demonstrate the benefits of WAVES, we run multi-threaded applications from SPLASH-2 [11] and PARSEC [12] benchmark suites. For each experiment, we execute 10 billion instructions in the region of interest. Our baseline policy activates all laser wavelengths ($\lambda_{act} = \lambda_{tot}$). We experiment with three different L_{thr} values and activate minimum laser wavelengths ($\lambda_{act} = \lambda_{min}$). Figure 3 shows the overall PNoC power savings with WAVES.

We observe that larger thread counts result in increased inter-chiplet network traffic among the communicating threads and, therefore, desire higher λ_{min}. In addition, network-intensive applications such as *swaptions, cholesky* and *canneal* desire higher λ_{min} than other non-network-intensive applications, and therefore, result in lower power savings. As the on-chip thermal gradient increases due to higher logic power, thermal tuning power increases resulting in lower power savings in such network-intensive applications. Overall, we obtain 23% (resp. 38%, 42%) average PNoC power savings with only 1% (resp. 5%, 10%) performance loss.

IV. Conclusion

PNoCs are developing as promising alternatives to ENoCs for addressing the communication energy bottlenecks in large manycore systems. However, it is critical to address the high PNoC power consumption arising from laser, EOE and MRR thermal tuning. We present a low-latency and low-storage wavelength selection technique, WAVES, that accounts for on-chip TV and PV and activates the best λ_{min} at runtime, without degrading the system performance.

Acknowledgement

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] 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, pp. 1246–1260, 2008.

[2] A. Joshi *et al.*, "Silicon-photonic clos networks for global on-chip communication," in *Proc. NoCS*, 2009, pp. 124–133.

[3] W. Bogaerts *et al.*, "Silicon microring resonators," *Laser & Photonics Reviews*, vol. 6, no. 1, pp. 47–73, 2012.

[4] J. Cardenas *et al.*, "Low loss etchless silicon photonic waveguides," *Optics express*, vol. 17, no. 6, pp. 4752–4757, 2009.

[5] Y. Thonnart *et al.*, "A 10Gb/s Si-photonic transceiver with 150μW 120μs-lock-time digitally supervised analog microring wavelength stabilization for $1Tb/s/mm^2$ die-to-die optical networks," in *Proc. ISSCC*, 2018, pp. 350–352.

[6] J. Howard *et al.*, "A 48-core IA-32 message-passing processor with DVFS in 45nm CMOS," in *Proc. ISSCC*, 2010, pp. 108–109.

[7] A. Narayan *et al.*, "WAVES: Wavelength Selection for Power-Efficient 2.5D-Integrated Photonic NoCs," in *Proc. DATE*, 2019.

[8] 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, p. 52.

[9] S. Li *et al.*, "McPAT: an integrated power, area, and timing modeling framework for multicore and manycore architectures," in *Proc. MICRO*, 2009, pp. 469–480.

[10] K. Skadron *et al.*, "Temperature-aware microarchitecture," in *Proc. ISCA*, 2003, pp. 2–13.

[11] S. C. Woo *et al.*, "The SPLASH-2 programs: Characterization and methodological considerations," in *ACM SIGARCH computer architecture news*, vol. 23, no. 2, 1995, pp. 24–36.

[12] C. Bienia *et al.*, "The PARSEC benchmark suite: Characterization and architectural implications," in *Proc. PACT*, 2008, pp. 72–81.