

Non-Volatile Phase Change Material based Nanophotonic Interconnect

¹Parya Zolfaghari, ²Joel Ortiz, ²Cedric Killian, ¹Sébastien Le Beux

¹Department of Electrical & Computer Engineering, Concordia University, Montreal, Canada

²Univ Rennes 1, Inria, CNRS/IRISA, Lannion, France

p_zolfag@encs.concordia.ca, joel.ortiz-sosa@inria.fr, cedric.killian@irisa.fr, slebeux@encs.concordia.ca

Abstract— Integrated optics is a promising technology to take advantage of light propagation for high throughput chip-scale interconnects in many core architectures. A key challenge for the deployment of nanophotonic interconnects is their high static power, which is induced by signal losses and devices calibration. To tackle this challenge, we propose to use Phase Change Material (PCM) to configure optical paths between writers and readers. The non-volatility of PCM elements and the high contrast between crystalline and amorphous phase states allow to bypass unused readers, thus reducing losses and calibration requirements. We evaluate the efficiency of the proposed PCM-based interconnects using system level simulations carried out with SNIPER manycore simulator. For this purpose, we have modified the simulator to partition clusters according to executed applications. Simulation results show that bypassing readers using PCM leads up to 52% communication power saving.

Index Terms— Phase change Material, ONoC, Design Method

I. INTRODUCTION

Advances in fabrication process enabled the use of manycore architectures for High Performance Computing (HPC) [1]. Due to their increasing complexity and the heterogeneity of targeted applications, task mapping became critical to optimize execution time, energy efficiency and reliability [2]. Considering the large number of diversity of cores, memories, and accelerators, Network on Chips (NoCs) naturally became the backbones of manycore architectures. This further complexifies task allocation [3], which now require complex algorithms involving, for instance, neural networks [4]. On the performance side, nanophotonic interconnects, also called Optical Networks-on-Chip (ONoCs), have the potential to deliver the bandwidth required by data intensive applications in HPC [5]. However, they suffer from high static power consumption, which is mostly due to losses experienced by optical signals and devices calibration requirements. Indeed, optical devices are sensitive to manufacturing process and thermal variations, which call for constant calibration of resonating devices such as ring resonators [6].

To reduce the static power consumption of nanophotonic interconnects, we investigate the use of Phase Change Material (PCM) elements to bypass unused optical devices. Indeed, the high contrast between crystalline (cr) and amorphous (am) phase states has enabled the design of non-volatile directional couplers [7]. A key limitation of PCM elements is the slow phase state changes compared to the required nanosecond scale latency communication requirements in manycore. We tackle this challenge by partitioning the manycore to execute different applications and by reconfiguring the interconnect only when new applications are executed. While recent studies have

demonstrated the efficiency of PCMs to design multi-level optical memories [8] and weighting functions in spiking neural networks [9], this work is the first attempt to use the technology in nanophotonic interconnects.

The contributions are the following. We propose a non-volatile Single Writer Multiple Reader (SWMR) optical channel. The channel is reconfigurable and allows to bypass unused readers using PCM-based directional couplers. We develop a method to configure the interconnect according to the mapping of applications on the cores. To do this, we modified SNIPER simulator to enable the thread distribution of multi-applications from Splash2 and PARSEC benchmarks. We evaluate the power consumption of the network using a model we developed. We show that an average of 21% power saving is obtained compared to channels without PCM, and up 52% saving is reached for mapping involving 4 clusters. Finally, we simulate the execution of the applications in parallel by considering the proposed reconfigurable interconnect. This allows us to divide the architecture to execute two applications independently and simultaneously. On average, we obtain a 21.6% reduction in the execution time.

The paper is organized as follows. Section II discusses related works. In Section III, we present the proposed PCM based optical interconnect. Section IV describes the power model. Section V presents the experimental setups and Section VI is results section. Section VII concludes the work.

II. RELATED WORK

A. Optical Network on Chip

The static power consumption of nanophotonic interconnects is mostly due i) losses experienced by optical signals, which call for high laser power, and ii) optical devices calibration [10]. At the link level, PID controllers are designed to compensate the effect of temperature variation on the wavelength shift of ring resonator resonance [11]. While the design in [11] involves to monitor received optical signals, the controlled proposed in [12] relies on sensors monitoring the temperature of ring resonators. At the network level, reconfigurable interconnects were proposed to optimize the use of optical channels [5], thus allowing to reduce the resources overhead. At system level, application-specific mapping methods allow to reduce crosstalk noise [13]. Run-time methods enabling to select the number of used wavelengths according to bandwidth requirement was proposed in [6]. Approximate computing techniques were also investigated to reduce bandwidth requirement [14]. All these approaches are complementary to this work, which aims at using PCM to reduce the static power.

B. Phase Change Material (PCM)

PCM elements are characterized by femtojoule-scale phase transition energy consumption, 10^{15} switching cycle endurance and years long state retention [15]. These properties led to the usage of PCMs in numerous optical application such as spiking neural networks [9], optical memories [8] and directional couplers [7]. The design proposed in [7] relies on a PCM element inserted between the branches of a directional coupler. The signal transmission depends on the phase state of the material: amorphous (*am*) and crystalline (*cr*) states lead to cross and bar transmission respectively. This work involves such a directional coupler, which we use to bypass disconnected readers in SWMR channels.

III. PCM BASED SWMR CHANNEL

This section presents the proposed PCM-based nanophotonic interconnect. It involves SWMR links onto which PCM elements are added to control the routing of the optical signals. We first present an overview of the proposed interconnect and we then present use-case scenarios.

A. Overview

Figure 1 represents the proposed configurable SWMR channel assuming one waveguide and N wavelengths. As in conventional SWMR channel, the writer modulates the optical signal, which propagates towards the destination receiver where Opto-Electronic (OE) conversions are carried out. However, unlike conventional SWMR channels where all intermediate receivers must be crossed before reaching the destination, the configuration ability of the proposed SWMR channel allows to connect only a selected set of readers. By disconnecting readers from the optical path, we aim at reducing the power consumption as follows: i) MRRs in disconnected receivers do not require, power hungry, calibration and ii) optical losses can be reduced since signals do not propagate through all the MRRs. To achieve this, we introduce a bypass waveguide at each reader, and we use a Directional Coupler (DC) between each reader to configure the optical path as represented in Figure 1. For each DC, the path can be configured to transmit signals towards the bypass waveguide or towards the receiver. The couplers are large band, thus allowing transmitting all modulated signals to a same output waveguide. The DC is configured according to the state of an embedded PCM element, as shown in Figure 1.c:

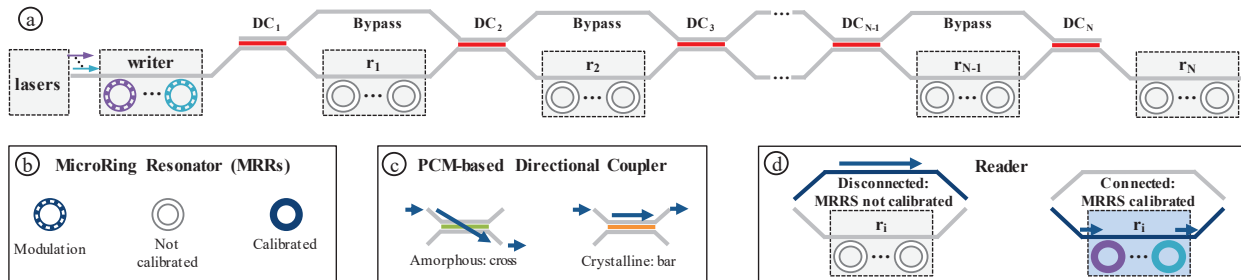


Figure 1: a) proposed SWMR channel with PCM-based directional couplers to configure the optical path through readers and bypass waveguides, b) MRRs states, c) signal transmission (for all wavelengths) through the directional coupler according to the state of the PCM and d) signal transmission to connected reader or use bypass path for disconnected reader.

amorphous (*am*) and crystalline (*cr*) phases lead to cross and bar configuration respectively. The PCM is non-volatile, thus allowing maintaining the DC configuration without consuming static energy.

B. Configuration Method and Use Case Scenarios

Considering the rather slow state change of PCM elements (around 100ns [16]) and their limited endurance, the SWMR channels are configured according to application specific or architecture specific connectivity requirements. The execution times of the targeted benchmarks applications, which typically range from 100ms to 10s, will ensure a low reconfiguration frequency of the SWMR channel ($\ll 1$ Hz). The connectivity requirements depend on the number of writers (i.e. the number of SWMR channels) and number of readers per channel. In the context of manycore architectures where multiple SWMR channels are used, each channel is configured according to the list of readers to be reached by the writer. To connect a reader, the preceding DC is configured either to cross if the previous reader is disconnected (i.e. signals coming from bypass) or to bar if the previous reader is also connected. Since we assume that a SWMR configuration does not change during the execution of an application, the MRRs of disconnected readers are not calibrated. On the contrary, the MRRs of connected readers are calibrated as follows: MRRs of the destination receiver drop the signals from the waveguides, while the others MRRs let the signals propagating (through). Maintaining the calibration of all connected MRRs allows minimizing the communication latency overhead [12]. To further save energy, the injected laser power is adapted to losses experienced by the optical signal for each configuration.

Figure 2 illustrates various connectivity scenarios for a 4-readers SWMR channel. We assume three wavelengths, thus leading to three MRRs per reader. The scenarios are:

- **all readers connected** (Figure 2.a) leads to a regular SWMR channel for which all the MRRs require calibration. To achieve the channel configurations, all the PCM elements are set to *cr* state, thus leading to bar transmission for the directional couplers. As previously mentioned, any of the connected receivers can be reached by the writer by calibrating the corresponding MRRs to the drop transmission .

(the others MRRs are calibrated to through transmission). Regarding the optical losses, this configuration leads to the highest losses, and hence highest laser power requirements, since the optical signal can propagate through all MRRs.

- **r₂, r₃ and r₄ connected** (Figure 2.b) require to bypass r₁, which is achieved by configuring DC₁ to the *am* state. Since r₂ is part of the connected readers, DC₂ is also configured to *am*. The rest of the PCMs are set to *cr* state, thus allowing to transmit the signals to r₃ and r₄. Since one reader is bypassed, the MRR through losses are reduced, thus leading to laser power saving. Furthermore, reduction of calibration power is also achieved since signals do not propagate through r₁.
- **r₁ and r₃ connected** (Figure 2.c) involves bypassing r₂. PCM element in DC₁ is configured to *am*, thus allowing the signal propagation through r₁. To bypass r₂, PCM elements in DC₂ and DC₃ are also configured to *am*. Since r₄ is also disconnected and is beyond the last connected reader, DC₄ does not require any specific configuration. In term of power, further reduction can be achieved compared to previous scenario due to lower MRRs losses, waveguide propagation losses and MRR calibration.
- **r₂ connected** (Figure 2.d) leads to SWSR channel. To achieve this, DC₁ and DC₂ are set to cross state while DC₃ and DC₄ can be in any state (i.e. no reconfiguration needed).




Figure 2. PCM elements configuration and ring calibration for various scenario: a) all interfaces connected, leading to a regular SWMR channel, b) r₁ disconnected, which allows to not calibrate its MRRs, c) r₂ and r₄ disconnected d) r₂ only connected leading to SWSR channel.

IV. POWER MODEL

In this section, we first describe the method to configure the SWMR channels according to connectivity requirements, then we present the network power model.

A. Method

We estimate the total network power consumption by considering the contribution of each SWMR channel, as defined in Eq. 1. P_{SWMR_i} corresponds to the power consumption of channel *i*, which depends on its configuration. In Eq. 2, P_{laser_i} is the laser power required to reach the photo detector sensitivity on the last connected reader. P_{MR-C_i} is the power required to calibrate the rings in connected readers and P_{config_i} is the power needed to configure the PCM elements in the directional couplers. P_T is the power required to serialize and modulate optical signals in the writer. P_R is the receiver power

consumption, which corresponds to photo detection, amplification, comparison and deserialization.

$$P_{total} = \sum_{i=0}^{N-1} P_{SWMR_i} \quad (1)$$

$$P_{SWMR_i} = P_{laser_i} + P_T + P_R + P_{config_i} + P_{MR-C_i} \quad (2)$$

To evaluate laser, calibration and reconfiguration powers, we consider the following key parameters from the connectivity requirement: i) number of connected readers *n*, ii) number of PCM elements in the crystalline state *A*, iii) number of PCM elements in the amorphous state *B* and iv) position of the last connected reader p_{r_i} .

B. Laser Power

For each channel, we define the required laser power according to the photodetector sensitivity (typically -8 dBm for 10⁻¹² BER [14]), the lasing efficiency and the worst-case losses experienced by the signal to reach the last connected reader (Eq. 3). The worst-case losses in a channel is defined by Eq. 4

$$P_{laser_i} = (P_{received_i} + Loss_{wci})/eff \quad (3)$$

$$Loss_{wci} = L_{MR-t} \times N_{\omega} \times n + L_w \times p_{r_i} \times d +$$

$$L_{MR-d} + \sum_{a=0}^{A-1} IL_{cr}^{bar} + \sum_{b=0}^{B-1} IL_{am}^{cross} + X_{talk} \quad (4)$$

The propagation of the signals towards the last connected reader involves crossing all rings in the *n* connected interfaces. L_{MR-t} is the through loss per ring and N_{ω} is the number of wavelengths, which also corresponds to the number of rings per reader. The waveguide propagation loss experienced by the signals depends on the position of the last connected reader, the distance between interfaces (*d*) and the waveguide loss (L_w in dB/cm). L_{MR-d} is the drop loss in the receiver. The insertion loss induced by the directional couplers depends on the number of crossed PCM elements and their state. X_{talk} is crosstalk power penalty which depends on the number of wavelengths and the ring position in the reader interface [17].

C. Ring Calibration

MRRs are sensitive to temperature variation, which calls for calibration to compensate the wavelength detuning. To evaluate the calibration power, we consider the Thermal Tuning Efficiency (TT_e) of the closed loop feedback system proposed in [6]. Only the rings in connected readers require calibration, which leads to Eq. 5. In the equation, $\Delta\lambda_{MR-C_{kj}}$ corresponds to the required wavelength shift of ring at position *j* in reader *k*. Assuming a homogeneous distribution of the signal wavelengths among the FSR, the maximum shift for each ring is FSR/N_{λ} . Therefore, the required shift for each MRR is defined by Eq. 6 where $\Delta\lambda_{shift_{kj}}$ is the wavelength shift of MRR *j* located at reader *k*. It is obtained with Eq. 7 where ΔT is the temperature variation and $\frac{d\lambda}{dT}$ is the MRR thermal sensitivity.

$$P_{MR-C} = \sum_{k=1}^n \sum_{j=1}^{N_{\lambda}} \Delta\lambda_{MR-C_{kj}} * TT_e \quad (5)$$

$$\Delta\lambda_{MR-C_{kj}} = \frac{FSR}{N_{\lambda}} - (\Delta\lambda_{shift_{kj}} \bmod \frac{FSR}{N_{\lambda}}) \quad (6)$$

$$\Delta\lambda_{shift} = \frac{d\lambda}{dT} \Delta T \quad (7)$$

$$P_{config} = f \times (\sum_{a=0}^{N_{PCM}^{cr}} E_{cr_i \rightarrow am_i} + \sum_{b=0}^{N_{PCM}^{am}} E_{am_j \rightarrow cr_j}) \quad (8)$$

D. Channel Configuration

The following presents a model to evaluate the power needed to reconfigure SWMR channels. Indeed, when executing a new set of applications, the interface connectivity might change, thus leading to a different set of connected readers. This requires changing the state of PCMs from *cr* to *am* or from *am* to *cr*, which requires energy. By assuming PCM already configured for a given connectivity requirement, we obtain reconfiguration power using Eq. 8. It is worth mentioning that the reconfiguration time of the PCMs is not critical in our system since SWMR channels are reconfigured only when new applications are executed. Hence, the design of circuits to change PCM state is out of the scope of the paper but details related to implementation involving electrical pulses, optical pulses or external heaters can be found in [18].

V. EXPERIMENTAL SETUP

A. Architecture

We consider the 3D architecture illustrated in Figure 3: bottom layer implements processing units and memories, while top layer implements the optical interconnect. It is composed of 16 clusters of 4 cores. Each cluster includes a shared last level cache (L3). Each core has a private L1 data (L1d), instructions (L1i) caches, of 32KB each, and a private L2 cache of 512KB. MESI protocol ensures coherency between the distributed caches. All clusters are connected with ONIs through TSVs. Each ONI has 1 transmitting and 15 receiving waveguides which are featured with SWMR and WDM, and include PCM based bypassing paths, as depicted in Figure 3. Each waveguide transmits 8 optical signals at different wavelengths.




Figure 3. Considered 3D hardware architecture.

B. Simulator and Benchmarks

We consider PARSEC and SPLASH2 benchmark suites [20] which are commonly used to study the performance of manycore architectures. They contain applications from different domains, which leads to various processing and communication requirements. We have selected six representative applications according to their parallelization level: i) FFT and Raytrace from Splash2 and ii) x264, Blackscholes, Barnes and Dedup from PARSEC. The execution of these applications on the architecture are simulated using the SNIPER manycore simulator [19]. We modified the Splash2 applications to allow their parallel executions and we modified SNIPER to handle the distribution of threads. It is possible to distribute applications on any number of clusters, which is suitable to explore the design space (see results section). Although the application mapping and the cluster partitioning lead to numerous design options, the results will only refer to scenario below due to space limitation.

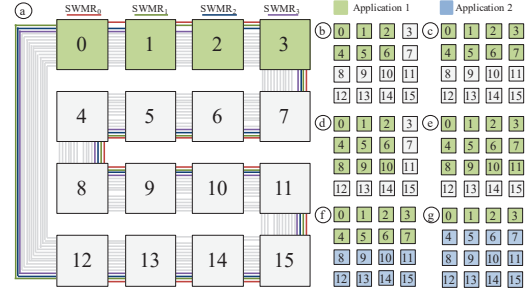


Figure 4: mapping example: a) application 1 mapped on Cluster 0 to 3, which leads to the use of Channel 0 to 3 only ; b-e) mapping of one application; f-g) mapping of two applications.

Figure 4.a illustrates the mappings of an application on four clusters, which leads to the use of four SWMR channels. Each used channel is configured to bypass disconnected readers (not illustrated in the figure for sake of clarity). Figure 4.b-e illustrate mappings on 6, 8, 9 and 12 clusters. Figure 4.f-g illustrate cluster partitioning for the parallel execution of two applications. From the simulation results, we extract the communication traces between L2 caches and the distributed L3, which corresponds to the traffic on the SWMR channels. Table 1 summarizes the architectural and technological parameters.

Table 1: Hardware and Technological Parameters.

Parameter	Definition	Value
Architecture		
# clusters		16
# cores per cluster		4
L1/D cache		32 KB each
L2 cache		512 KB
Cache protocol		MESI
ONI		
d	Distance between interfaces	0.376cm [14]
N_ω	Number of wavelengths	8
B_R	Bit-rate	10Gb/s
P_T	Transmitter power	24mw [6]
P_R	Receiver power	24mw [6]
L_w	Waveguide loss in db/cm	0.25 [14]
L_{MR-d}	MRR drop loss	0.7 [14]
L_{MR-t}	MRR through loss	0.02 [14]
Q	MRR Quality Factor	20,000 [17]
TT_e	Thermal tuning efficiency	120pm/mw [6]
X_{talk}	crosstalk power penalty	0.0494 dB
Directional Coupler		
IL_{cr}^{bar}	Bar output Insertion loss for PCM in cr state	0.16dB [7]
IL_{am}^{cross}	cross output Insertion loss for PCM in am state	0.72dB [7]
$E_{cr \rightarrow am}$	Phase transition energy from cr to am	2nj [7]
$E_{am \rightarrow cr}$	Phase transition energy from am to cr	2nj [7]

VI. RESULTS

A. Losses and Power Analysis

We investigate the power consumption of the proposed interconnect and we compare it with a baseline interconnect without PCM elements. For this purpose, we first evaluate the loss of each channel used for 1x4 SWMR configuration and we report results in Figure 5. For SWMR₀, the waveguide loss is as small as 0.28 dB due to the short distance between the writer (cluster 0) and the last connected reader (cluster 3). Since readers located after the last connected reader are not used, the MR through loss for ONoC with and without PCM is the same

(1.41dB). The need to cross three directional couplers for the PCM-based channel leads to 0.48dB overhead compared to the channel without PCM. Channels 1 to 3 demonstrate higher waveguide propagation losses since the signals propagate through the entire waveguide. Since only three readers are connected, MR through loss for channels with PCM is the same as for SWMR₀ (i.e. 1.41dB) while they reach 2.4 dB without PCM due to the need to cross all the MRRs in the intermediate readers. Finally, the losses induced by the directional couplers (3.56dB) lead to 1.4dB additional losses for the PCM based channel compared to channel without PCMs. To summarize, although PCMs allow reducing the MRR through loss when bypassing disconnected readers, the rather high insertion loss they involve lead to higher total losses. While PCM-based interconnect requires higher laser power, we will show in the following that significant power saving can be achieved thanks to significant reduction in the ring calibration power.




Figure 5: Loss breakdown on each channel for 1x4 configuration.

Figure 6.a shows the power breakdown for previously discussed configuration and channels. For SWMR₀, the power consumption is 1mW higher with PCM due to laser power overhead. Calibration power is the same with and without PCM since no reader is bypassed. For channels 1, 2 and 3, we obtain x5 reduction in the calibration power since only the connected readers require calibration. Overall, despite the 8mW increase in laser power, our approach leads to up to 52% power saving. As shown in .b, we obtain an average of 45% power saving per used channel power for 1x4 configuration. As the number of connected readers increases, lower power saving is obtained due to ring calibration requirements. Although configuration 4x4 leads to 6% power increase, our approach demonstrates a 21% power reduction on average. The higher energy efficiency of our approach for lower connectivity scenario ideally complements with the mapping of independent applications on the clusters, which we study in the following.




Figure 6: Power consumption results: a) power breakdown per channel for 1x4 configuration and b) average power consumption per channel according to the network configurations.

B. Benchmark Analysis

We simulate the execution of benchmark applications using the environment defined in Section V.B, as reported in Figure 7. The mapping of Blackscholes (Figure 7.a) on 4 clusters (resp. 16 clusters) leads to 1376ms (resp. 1188ms) execution time and 28W (resp. 120W) power consumption. By using all clusters in the architecture, a 15% speedup is achieved at the cost of x4.2 power increase. Configuration involving 4 clusters and 8 clusters lead to intermediate speedups and powers, which we will use to optimize the multi-applications mapping and the ONoC configuration. While all tested applications follow a similar trend, the actual speedup and power consumption highly depends on the parallelization level and the communication patterns. Using 4 clusters for FFT is preferable since 16 clusters lead to 3% speedup and 3.5x power overhead (Figure 7.e). On the other hand, the mapping of x264 on 16 clusters allows 60% speedup. To conclude, clusters partitioning, which is achieved by reconfiguring SWMR channels, can be optimized depending on the executed applications.




Figure 7: Execution time and power wrt. number of used clusters.

C. Multi-Application Mapping and ONoC Configuration

We consider the execution of pairs of applications on the architecture. The baseline is a sequential execution of the two applications on all the clusters (denoted 16→16 in Figure 8) connected using an ONoC without PCM. As reported in Figure 8.b, the baseline execution time for x264 and Blackscholes is 3663ms. We now consider the proposed PCM-based ONoC. PCMs allow to partition the architecture in order to allocate a dedicated number of clusters per application (see Figure 4.f for mapping details). Since readers which are not part of a SWMR channel partition are disconnected, the application tasks can be executed in parallel without any resources sharing, which contribute to reduce the execution time. As illustrated in Figure 8.b (case denoted 8/8 in the figure), the execution time is significantly reduced and reaches 2962ms. We then consider an uneven mapping of the application on the clusters (illustrated in Figure 4.g): 4/12 (resp. 12/4) allocates 4 (resp. 12) clusters to Blackscholes (resp. x264) and 12 (resp. 4) clusters to x264 (resp. Blackscholes). Mapping 12/4 leads to the best improvement in the execution time (26.8%). For the interconnect, SWMR channels are configured to connect only readers of interface executing the same application as the writer. Considering the 1.2Tb/s aggregated bandwidth of the optical interconnect (16 channels with 8 wavelengths at 10Gbps each), the baseline scenario would lead to a 2.7pJ/bit energy consumption under the assumption that all channels permanently transmit data. By considering the actual number of bits transmitted, we obtain a 47.7% energy per-bit reduction for the proposed ONoC compared to a network without PCM. The significant

improvement in energy is due to i) the reduced static power induced by bypassed readers and ii) the higher data rate induced by shortened execution time. We carried out the same study for Blacksholes/Raytrace application pair, as reported in Figure 8.c. Interestingly, the best configuration is obtained by allocating 4 clusters to Blacksholes (against 12 in previous case), which is due to a lower computation load of Raytrace. Mapping 4/12 leads to 40.3% execution time reduction and 42.4% energy per bit reduction.

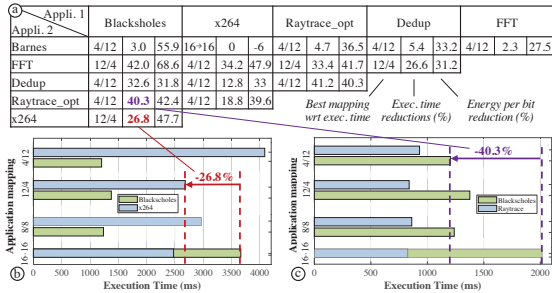


Figure 8: a) Best mapping results and improvements compared to execution on 16 clusters, b-c) execution time for different mappings and ONoC configurations for Blacksholes/x264 and Blacksholes/Raytraces.

Figure 8.a summarizes the best mapping we obtain for each application pair. Reductions in execution time and energy per bit reach up to 42% and 68.8% respectively for FFT/Blacksholes. Mappings involving Barnes lead to the lowest improvement since it is the least computation intensive application. No improvement is obtained when Barnes is combined with x264, which is the most data intensive application. This calls for finer grain mapping we will explore in our future works. On average, we obtain 21.6% execution time reduction and 41.2% per-bit energy reduction for the PCM-based ONoC. This validates the potential for PCMs to improve the on-chip optical communication energy efficiency and to facilitate partitioning of manycore.

In our last study, we evaluate the power required to change the state of PCMs during network reconfigurations. For this purpose, we first consider a pessimistic scenario in which the two applications with the smallest execution time, i.e. Barnes and Dedup, require full reconfiguration of the PCMs after each execution. The resulting 1.3Hz reconfiguration frequency would lead to 5μW, which is negligible considering the total system power consumption (100W). Nevertheless, the PCMs reconfiguration power and endurance can be improved by optimizing the remapping to take into account the current state of PCMs. This will be explored in our future works.

VII. CONCLUSION

This paper investigates the potential of PCMs elements for reducing static power in nanophotonic interconnects. We propose a configurable non-volatile SWMR channel capable of bypassing unused readers, thus avoiding unnecessary calibration of optical devices. To efficiently use the non-volatile capability of our interconnect, we develop a method to map applications on dedicated clusters. Configurability property of PCM allows to partition the architecture for parallel execution

of applications which leads to reduction in execution time compared with sequential execution of applications on all 16 clusters. Results show that up to 52% static power reduction of the interconnect and 42% execution time reduction can be reached. Since this work is the first attempt to use PCM elements in nanophotonic interconnects, the perspectives are numerous and include the bypassing of group of readers to reduce insertion losses and the fine grain mapping of applications on manycore.

ACKNOWLEDGMENT

This work benefited from the Fonds de recherche du Québec – Nature et technologies (FRQNT) under grant 2021-NC-286424 and the support of the SHNoC ANR18-CE25-0006 and OpticALL2 ANR-18-CE24-0027 projects of the French National Research Agency (ANR).

REFERENCES

- [1] P. Vivet et al., "IntAct: A 96-Core Processor With Six Chiplets 3D-Stacked on an Active Interposer With Distributed Interconnects and Integrated Power Management," IEEE Journal of Solid-State Circuits, 2021
- [2] A. K. Singh et al., "Mapping on multi/many-core systems: Survey of current and emerging trends," ACM/EDAC/IEEE Design Automation Conference (DAC), 2013
- [3] H. Chen et al., "MARCO: A High-performance Task Mapping and Routing Co-optimization Framework for Point-to-Point NoC-based Heterogeneous Computing Systems," ACM TECS, Sept. 2021
- [4] M. F. Reza et al., "Neuro-NoC: Energy Optimization in Heterogeneous Many-Core NoC using Neural Networks in Dark Silicon Era," IEEE ISCAS, 2018
- [5] X. Wu et al., "SUOR: Sectioned Unidirectional Optical Ring for Chip Multiprocessor," JETC, Vol. 10, pp 1–25, 2014
- [6] A. Narayan et al., "PROWAVES: Proactive Runtime Wavelength Selection for Energy-efficient Photonic NoCs", IEEE TCAD, 2020
- [7] P. Xu et al., "Low-Loss and Broadband Nonvolatile Phase-Change Directional Coupler Switches," ACS Photonics, 6(2), 2019
- [8] C. Ríos, M. Stegmaier, P. Hosseini, et al. "Integrated all-photonic non-volatile multi-level memory," Nature Photon 9, 725–732, 2015
- [9] I. Chakraborty, G. Saha, A. Sengupta, et al. "Toward Fast Neural Computing using All-Photonic Phase Change Spiking Neurons," Sci Rep 8, 12980, 2018
- [10] Chen Sun et al., "DSENT - A Tool Connecting Emerging Photonics with Electronics for Opto-Electronic Networks-on-Chip Modeling", IEEE/ACM NOCS, 2012
- [11] D. Dang et al., "PID controlled thermal management in photonic network-on-chip," 33rd IEEE ICCD, 2015
- [12] X. Chen et al. "Simultaneously Tolerate Thermal and Process Variations Through Indirect Feedback Tuning for Silicon Photonic Networks," IEEE TCAD, 2021
- [13] E. Fusella and A. Cilaro, "Crosstalk-Aware Automated Mapping for Optical Networks-on-Chip", TECS, Vol. 16, pp 1-26, 2016
- [14] J. Lee et al., "Approximate nanophotonic interconnects". IEEE/ACM NOCS 2019
- [15] J. Zheng et al., "GST-on-silicon hybrid nanophotonic integrated circuits: a non-volatile quasi-continuously reprogrammable platform," Opt. Mater. Express, 2018
- [16] N. Yamada et al., "Rapid, phase transitions of GeTe-Sb2Te3 pseudobinary amorphous thin films for an optical disk memory". Journal Applied Physics, 1991
- [17] M. Bahadori et al., "Crosstalk penalty in microring-based silicon photonic interconnect systems," J. Lightwave Technol., 2016
- [18] S. Abdollahramezani et al., "Tunable nanophotonics enabled by chalcogenide phase-change materials", Nanophotonics, 2020
- [19] T. E. Carlson et al., "An Evaluation of High-Level Mechanistic Core Models," ACM TACO, 2014
- [20] C. Bienia, et al., "PARSEC vs. SPLASH-2: A quantitative comparison of two multithreaded benchmark suites on Chip-Multiprocessors," IEEE International Symposium on Workload Characterization, 2008