HAL
open science

# Distance-aware Approximate Nanophotonic Interconnect

Jaechul Lee, Cédric Killian, Sebastien Le Beux, Daniel Chillet

## To cite this version:

Jaechul Lee, Cédric Killian, Sebastien Le Beux, Daniel Chillet. Distance-aware Approximate Nanophotonic Interconnect. ACM Transactions on Design Automation of Electronic Systems, 2022, 27 (2), pp.1-30. 10.1145/3484309 . hal-03500153

**HAL Id: hal-03500153**

**https://inria.hal.science/hal-03500153v1**

Submitted on 7 Dec 2022

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

# Distance Aware Approximate Nanophotonic Interconnect

JAECHUL LEE, Univ Rennes, Inria, CNRS, IRISA, France
CÉDRIC KILLIAN, Univ Rennes, Inria, CNRS, IRISA, France
SEBASTIEN LE BEUX, Concordia University, Canada
DANIEL CHILLET, Univ Rennes, Inria, CNRS, IRISA, France

The energy consumption of manycore architectures is dominated by data movement, which calls for energy-efficient and high-bandwidth interconnects. To overcome the bandwidth limitation of electrical interconnects, integrated optics appear as a promising technology. However, it suffers from high power overhead related to low laser efficiency, which calls for the use of techniques and methods to improve its energy costs. Besides, approximate computing is emerging as an efficient method to reduce energy consumption and improve execution speed of embedded computing systems. It relies on allowing accuracy reduction on data at the cost of tolerable application output error. In this context, the work presented in this paper exploits both features by defining approximate communications for error tolerant applications. We propose a method to design realistic and scalable nanophotonic interconnect supporting approximate data transmission and power adaption, according to the communication distance to improve the energy efficiency. For this purpose, the data can be sent by mixing low optical power signal and truncation for the Least Significant Bits (LSB) of the floating-point numbers, while the overall power is adapted, according to the communication distance. We define two ranges of communications, short and long, which require only four power levels. This reduces area and power overhead to control the laser output power. A transmission model allows estimating the laser power according to the targeted BER and the number of truncated bits, while the optical network interface allows configuring, at run-time, the number of approximated and truncated bits and the laser output powers. We explore the energy efficiency provided by each communication scheme and we investigate the error resilience of the benchmarks over several approximation and truncation schemes. The simulation results of ApproxBench applications show that, compared to an interconnect involving only robust communications, approximations in the optical transmission led up to 53% laser power reduction with a limited degradation at the application level with less than 9% of output error. Finally, we show that our solution is scalable and lead to 10% reduction in the total energy consumption, 35x reduction in the laser driver size and 10x reduction in the laser controller compared to state of the art solution.

CCS Concepts: • **Computer systems organization** → **Embedded systems**; *Redundancy*; Robotics; • **Networks** → Network reliability.

Additional Key Words and Phrases: Photonic on silicon, network-on-chip, approximate communications, manycore architectures

**ACM Reference Format:**
Jaechul LEE, Cédric KILLIAN, Sebastien LE BEUX, and Daniel CHILLET. 2021. Distance Aware Approximate Nanophotonic Interconnect. 1, 1 (December 2021), 30 pages. https://doi.org/10.1145/1122445.1122456

Authors' addresses: Jaechul LEE, Univ Rennes, Inria, CNRS, IRISA, F-22305, Lannion, France, jaechul.lee@irisa.fr; Cédric KILLIAN, Univ Rennes, Inria, CNRS, IRISA, F-22305, Lannion, France, cedric.killian@irisa.fr; Sebastien LE BEUX, Concordia University, Dept. of Electrical and Computer Engineering, Montreal, Quebec, Canada, slebeux@encs.concordia.ca; Daniel CHILLET, Univ Rennes, Inria, CNRS, IRISA, F-22305, Lannion, France, daniel.chillet@irisa.fr.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
© 2021 Association for Computing Machinery.
Manuscript submitted to ACM

## 1 INTRODUCTION

Nowadays, large-scale applications such as scientific computing, video/image processing and data mining require high performance with large amount of computation and data exchange communication. Manycore architectures, which integrate hundreds or thousands of cores within a same chip, fulfill the needs of hardware performance with huge parallel computation capacities. In these architectures, the on-chip interconnects are a key element to ensure communication needs, which directly affects the application performance, either on energy and execution time.

In the last decade, Electrical Network-on-Chip (NoC) appears as an efficient solution for multicore architectures, in the range of tens of cores on a-chip, to circumvent the parallelism limitations of traditional buses. Nevertheless, as manycore era progress, electrical NoCs suffer from scalability in terms of latency and energy due to a huge increase on the number of hops between cores [1, 2]. Optical NoC (ONoC), which bases on Nanophotonic interconnects, is a promising solution to overcome bandwidth and latency issues, as optical signals propagate near speed-of-light in waveguides. However, their implementations remain challenging due to the low efficiency of the lasers, which are key devices in such interconnects. Indeed, the laser power consumption is mainly driven by the high signal power required to transmit data at low Bit Error Rate (BER) for accurate communications [3].

Energy efficiency of ONoC is addressed at several aspects. At circuit level, in [4] proposes to reduce the power loss which driven from components such as reducing the number of Micro Ring resonators (MR) and waveguide crossings. At data level, the use of error correcting codes has been explored to allow transmission power reduction, leading to transmission errors that are compensated by data redundancies [5]. At communications level, communication scheduling can improve the energy efficiency thanks to less parallel communication leading to crosstalk and energy penalty [6–8].

Beside these on-chip interconnect evolution, approximate computing is an emerging solution to improve energy efficiency and execution speed of embedded computing systems [9, 10]. It relies on accuracy reduction of the data representations, which allows lowering design constraints and improves performances at the cost of Quality of Result (QoR) degradation measured with output errors [11]. It has been deployed in numerous manners on interconnects [12, 13], operators and memory levels [14, 15] as well. With various methods of approximate computing implementation, it offers storage efficiency, computation speed and power consumption improvements, while keeping a tolerable degradation at the application side.

In this paper, we propose a distance aware approximate nanophotonic interconnect to improve power consumption of on-chip communication. Our proposal classifies the communications by the distance between the source and destination and its data type for approximate communication. Only two distances of communication are considered: short and long-distance range. Based on this classification, laser power levels are adapted according to the application's error tolerance. For this latter, we propose to combine the use of approximation and truncation of data. Approximation allows to lower the laser power level, hence to save energy, by targeting a given quality of transmission, accordingly with the application output errors tolerance; while truncation is the most efficient way to save power as bits are not sent, hence some lasers are kept off, but at the cost of more errors at the application level. The proposed solution is a straightforward method to enhance the power efficiency of ONoC, as it only requires 3 power levels at each source.

The contributions of the paper are as follows:

(1) The association of approximation, truncation and Short/Long distance range management is presented and generalized. We also introduce the Optical Network Interface that allows the use of the proposed method with several examples that illustrate the concept details;

(2) A graphical method allows to determine the three power levels required in function of the targeted BER and to determine the trigger for short and long-distance range;
(3) A comprehensive study is performed on real life benchmarks to analyze the impact of approximation and truncation on the quality of results (QoR). This study allows to show that not all the applications are good candidates for these kind of low power techniques, and also highlights that a same approximation and truncation configuration can be applied to several applications;
(4) An analysis of power reduction is provided to highlight the contribution of each low power techniques;
(5) An exploration of the design space is done to show the power reduction versus quality of result. This study allows comparing the proposed technique with state-of-the-art ones and with the optimal solutions.
(6) An analysis of the driver laser complexity is performed which allows a comparison of our lightweight solution and more complex solution from state-of-the-art one.

Our contribution is validated through simulation with SNIPER manycore simulator [16], which carried out Approx-bench [11] benchmark applications running on shared-memory architectures. The result of evaluation for the proposal shows drastic laser power reduction of 53% for Streamcluster application at the cost of less than 8% of errors. Finally, we show that our solution is scalable and outperform state of the art solution as follows: 10% reduction in the total energy consumption, 35x reduction in the laser driver size and 10x reduction in the laser controller.

The rest of the paper is organized as follows. Section 2 introduces state of art of ONoC and approximate computing concept. Section 3 presents the proposed technique and its management, according to the application needs. Section 4 presents the considered architecture implementation and the simulation setup, and a traffic analysis of the considered application benchmarks. Section 5 provided a comprehensive analysis of the impact of approximation and truncation on the output errors of applications. In Section 6, we first explore the impact of both system and device level on the laser power. Then, it shows the power improvement obtained according to the application error. Finally, we provide a comprehensive study on the design complexity of the drivers for our method and [13]. Section 7 provides discussions and perspectives on the deployment of the approach to off-chip lasers-based interconnects, and on task mapping optimization opportunities Section 8 concludes this work.

## 2 RELATED WORK

On-chip nanophotonics interconnect is a promising replacing technology for electronic on-chip interconnect due to its high bandwidth and low latency. However, active optical devices require high power and they suffer power hungry. To overcome this problem, several attempts propose to improve power efficiency on optical interconnect. In [4, 17], an all optic interconnect topology suitable for managing optical communications is proposed. It is based on routers that reduce number of wavelengths and other optical components to reduce power consumption. This solution is compared to other existing topologies. In [18], the authors propose photonic on-chip network architecture composed of several sub-nets. This allows high performance and low latency by sharing optical channel between sub-nets and by dynamically schedule them without degrading throughput. Error correction code (ECC) is introduced in Optical NoC to reduce the power consumption [5]. The transmitter includes an ECC and the receiver detects faulty bits and can then correct them. Since additional faulty bits are acceptable using ECC, the optical power emitted by the laser source can be significantly reduced. In [8], the authors propose crosstalk aware wavelength scheduling. The correlation of application execution time and the energy consumption overhead is strategically analyzed to reduce the crosstalk induced by WDM (Wavelengths Division Multiplexing). In [6], the paper proposes dynamic bandwidth scaling based on network traffic

prediction. It dynamically tunes optical channel link according to the needs of bandwidth. By turning off the unused network link, it optimizes channel utilization and reduces optical power. The technique consists on configuring the optical light intensity based on the number of path used in a tree based optical NoC. In [19–21], based on network load of required by applications, the authors propose a network management. In [19], the authors propose to increase the laser power efficiency through the management of online activating and deactivating of L2 cache and the associated optical links. In [20], the off-chip laser power is dynamically managed in runtime to adapt the network to application characteristic, hence reducing laser power consumption. In [22] proposed heterogeneous multicore architecture consists of CPU-GPU cores. It reduced static optical power by predicting network load with link reservation. In, [23], the authors propose to optimize the proposed runtime power management using token based time division multiplexing. It switches ON/OFF laser source according to NoC bandwidth needs. In [24], the authors propose to balance the power consumption of lasers and the average packet latency on the ONoC. This is done by optimization of the photonic link parameters at design time, and by on-line trade-off management of power consumption and performances by limiting the cross-talk effects involved by the Micro-Ring Resonators.

In addition to these optical link optimizations, other publications propose to exploit data approximation. Indeed, approximate computing is touted as an efficient solution to reduce energy and execution speed of computing systems [9]. This paradigm has been integrated into electrical NoC architecture with various techniques that reduce the number of packets and decrease precision of data packet in NoC. In [25], the authors propose approximate framework for ENoC. The work presents data transmission approximation that reduces number of data transmission by transmitting approximately similar data. Their work implements specific routing header to recognize the data approximation configuration. The work presented in [26] proposes to regulate the packets injected in NoC to mitigate network congestion. Their proposal adoptively drops fraction of data packets that can be approximately predicted as a lost data. An approximate error correction control in NoC has been proposed for large-scale parallel systems in [27, 28]. This paper proposes to configure different lengths of error correction for the data packet. For that, it decreases the number of data bits that are under the error correction. They achieve latency and power efficiencies by packet prediction and approximate error checking technique on NoC. In [29], the authors propose approximate communications on the NoC through data compression and relaxation on synchronization. Their work improves the energy consumption and speedup performance by reducing data synchronization packet that causes network bottleneck. The paper [30] proposes dual voltage communication for approximate data communication in electrical NoC approximate communication by managing the power with 2 voltage levels that correspond to content of the packet. Headers and important data flits are accurately transferred at a high voltage while the approximatable flits are set to low voltage which may provoke bit flips at the destination. Overall, these approximate communications in NoC, the output quality management is necessary for application performance. The work presented in [31] addresses runtime management to controls the output quality of approximation. This paper investigates online quality management system for approximate computing. It proposes to detect the output data containing error coming from approximate computing, and to correct them to ensure the output quality of application.

These successful study cases of approximate computing in electrical NoC have been introduced into Optical interconnect NoC as well. On-chip optical interconnects communicate by emitting optical signals from the laser drivers into waveguides. The input power of laser is configured by estimating various device level characteristics (e.g. photodetector sensitivity, laser driver efficiency, waveguide loss, and microring resonator insertion loss) and system level parameters

(e.g. number of wavelengths, propagation distance). The optical laser driver which generates optical signal is a key elements and dominant power consuming device due to its efficiency [32]. To address this large power consumption from the laser driver and reduce the power consumed in laser source, various solutions have been proposed. To ensure accurate optical communications between source and destination cores, the power received by photodetectors must be higher than their sensitivity. If this received power is reduced, the transmission BER increases and this leads toward approximate communications. In our previous work [12], we introduced the first nanophotonic interconnect in which data approximations are performed to save power for error tolerant applications. It leverages error tolerant data with respect to its significance by transmitting them with lowered laser power. In [13], the authors propose to truncate bits by turning off lasers, thus allowing to save power. However, the truncation is only applied when the estimated power at the receiver is lower than the photodetector sensitivity, which does not allow exploring accuracy power tradeoffs. Furthermore, in [13], the authors propose loss-aware power configuration that adapts the laser power according to the communication distance between source and destination cores. However, the complexity of the laser driver, which has to be designed to numerous optical power levels, and the overhead of the run-time management are not evaluated. In order to achieve fine grain distance loss-aware for each optical link [13, 33, 34], it requires to have large amount of power level configurations and it requires large amount of effort in terms of computation and which may not be suitable for large size architectures.

In this work, we improve the optical interconnect based on a realistic and scalable laser power management. First, we propose improved approximate communication, and we investigate not only bit approximation but also bit truncation and association of these two techniques for each data communication. This allows going further on design exploration and achieving a point closest to optimal one. Thus it provides more power saving with very subtle output error control. This optimization of approximate configuration makes optical communication more robust and energy efficient while adapting error tolerance of application. Secondly, we propose a laser power management based on a technique taking into account the communication distance. Our proposal classifies the communications by distance between source and destination cores. However, unlike [13], only two distances are considered: short and long-distance ranges. This allows laser power configuration more suitable and realistic for various architectures in terms of design effort and energy efficiency. Third, this paper presents method of power level configuration and its computation for targeted BER as well as the determination of distance threshold. Throughout this design process, we could gain limited number of optical laser power levels while these levels still satisfy approximate communication and distance aware technique. For the last, this paper highlights all the different communication scenarios with approximate techniques and its power gain computation. According to these points, in results section, evaluation results are presented on several real life benchmarks, and a comparison with Lorax method is performed.

## 3 PROPOSED COMMUNICATION LINK

In this section, we first present the proposed interconnect architecture and the communication schemes supported. Then we detail the estimation of power saving obtained considering the communication distance.

### 3.1 Overview

Our proposal aims to manage, for a given application, power consumption regarding the Quality of Result (QoR). For this purpose, the proposed nanophotonic interconnect allows adapting the laser power level according to the communication distance and the transmitted data type. Without lack of generality, Fig. 1 illustrates our approach on a Single Writer

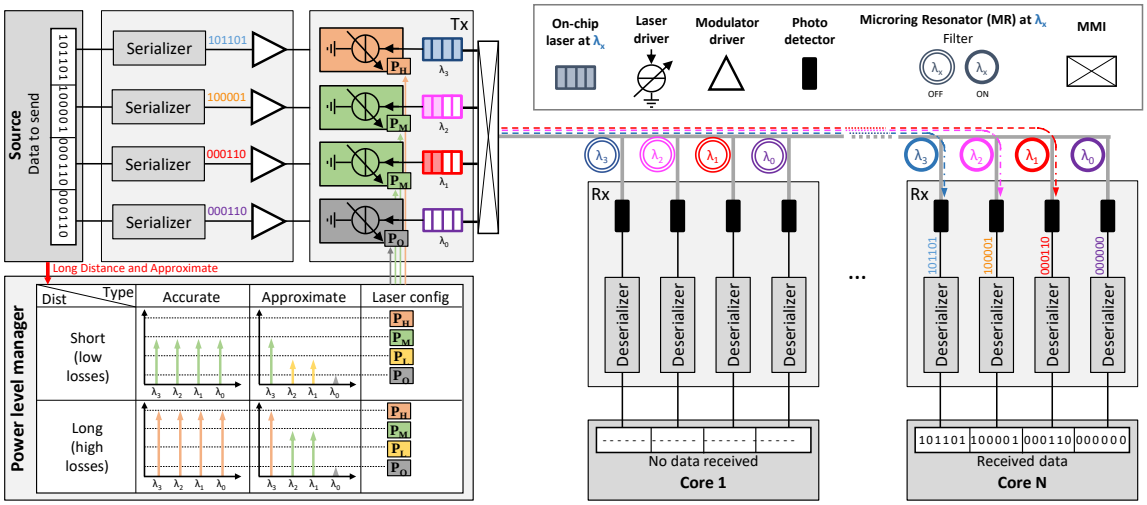


Fig. 1. Proposed data and distance aware optical link

Multiple Reader (SWMR) link. The writer (source) is connected to several readers (destinations represented by Core in the figure) using waveguides (for sake of clarity, only one waveguide is represented in the figure). Each waveguide allows transmitting $N\lambda$ signals using WDM ($\lambda_0..\lambda_{N\lambda-1}$). The signals are emitted by on-chip lasers and are combined into waveguide using MultiMode Interference (MMI) couplers [35].

When a communication is initiated, the writer activates the lasers (optical signals are emitted after few ns typically) and the data modulate the optical signals using On-Off Keying (OOK) modulation. Although we assume that data transmissions involve the use of all the optical signals, the transmission of data are serialized according to the data bit width ($N_{data}$) and the number of wavelengths ($N\lambda$). For instance, $N_{data} = 32$ and $N\lambda = 8$ would lead to a Bit Stream Length ($BSL$) of 4, i.e. 4 bits are sent on each wavelength and 4 clock cycles are needed to transmit each data. The design of the serializer has been investigated in [5, 36] and is out of the scope of the paper. The modulated signals propagate along the waveguide until reaching the destination. At this stage, the signals are ejected from the waveguide using MRs in ON state. They are redirected to photodetectors from which opto-electronic conversions are carried out. The serial streams are then de-serialized back into the original data format.

In case no communication occurs, the lasers are turned OFF and hence do not consume energy [20]. In case data transmission is triggered, lasers are turned ON with power levels that depend on i) the required communication quality and ii) the communication distance. To allow a reduction in consumption while limiting the driver complexity, we propose to use four laser power levels: high power ($P_H$), medium power ($P_M$) low power ($P_L$) and off ($P_{Off}$), with $P_H > P_M > P_L > P_{Off} = 0$. Fig. 1 presents an example assuming that the $\lambda_3$ is sent with high power, $\lambda_2$ and $\lambda_1$ are sent with medium power, while the power of $\lambda_0$ is equal to $P_{Off}$, meaning this wavelength is not used, which corresponds to 4 LSB bits truncation. The laser power values are calculated based on technological parameters (e.g. waveguide propagation loss and MR losses), system level parameters (e.g. number of readers) and the targeted BERs, as detailed in Section 3.3.

We assume that the accurate data (i.e. data that cannot be approximated) and approximated data require, on the receiver side, $BER_{NotAx}$ and $BER_{Ax}$ respectively. Hence, assuming $BER_{NotAx} < BER_{Ax}$, approximated data will be

transmitted using lower laser power compared to accurate data. As previously explained, the actual laser power depends on the propagation loss and hence the communication distance. For this purpose, we define two ranges of communication distance, *Short* and *Long*, depending on the position of the destination on the SWMR link. This leads to the four scenarios detailed in the following:

- **Accurate & Long** is used for the transmission of sensitive data (e.g. instructions and integer data) that critically impact the execution of applications. Therefore, the transmission quality is protected by targeting low BER for all the bits (e.g. $BER_{NotAx} = 10^{-12}$). To compensate the losses induced by the long-range transmission, the highest power level is used (i.e. $P_H$) for all the lasers.
- **Accurate & Short** is also used to transmit sensitive data but at a shorter distance. Due to the lower losses, low BER is obtained for all the bits by configuring the lasers power level to $P_M$.
- **Approximate & Long** allows the long range transmission of data that can be approximated. Depending on the approximation scheme, the lasers are configured with different power levels. As detailed in Section 3.2, most significant bits are transmitted with low BER, intermediate bits are transmitted with high BER and least significant bits are truncated. To achieve this, the lasers that emit signals transmitting most significant bits and approximated bits are configured to $P_H$ and $P_M$ respectively. Truncated bits are not transmitted and the corresponding lasers are thus set to $P_{Off}$.
- **Approximate & Short** allows for short distance transmission of data that can be approximated. As for the previous scenario, the configuration of the lasers depends on the bit's significance. However, the lower losses allow reducing the power of the emitted signals, which lead to use the $P_M$ and $P_L$ power level for most significant and intermediate bits respectively.

By adapting the laser powers according to both communication distance and level of approximation, our interconnect aims at minimizing the energy to transmit data bits. Since only three power levels are needed for all destinations with approximation levels, the design complexity of the laser driver remains comparable to solutions which require individual computation of power levels for each destination (*Distance Proportional*). The proposed method to define the laser power levels is presented in 3.3 and, in the following, we detail the considered data approximation scheme.

### 3.2 Data approximation scheme

The proposed communication method involves both approximation and truncation on data. In order to control different approximate level, we propose configurable bit areas according to the bit significance. Hence, several configurations are available to send data that can be approximated, and that can be tuned to achieve a given quality of result for an application. Fig 2 illustrates how the technique can be used for the IEEE 754-2008 single-precision Floating-Point (FP) format and for long distance communication. Bits are grouped into 3 different areas according to their significance on the floating-point data representation:

(1) $A_{NotAx} = [N_{data} - 1 \; ; \; P_{Ax}]$ corresponds to Most Significant Bits (MSB), composed with the sign and the exponent of the FP number, which are always protected from approximation thanks to a transmission power allowing to reach $BER_{NotAx}$;
(2) $A_{Ax} = [P_{Ax} - 1 \; ; \; P_{Trunc}]$ are intermediate bits, the highest part of the mantissa, which may tolerate approximation and are transmitted in order to reach $BER_{Ax}$;
(3) $A_{Trunc} = [P_{Trunc} - 1 \; ; \; 0]$ are the Least Significant Bits (LSB), lowest par of mantissa, and are not transmitted.

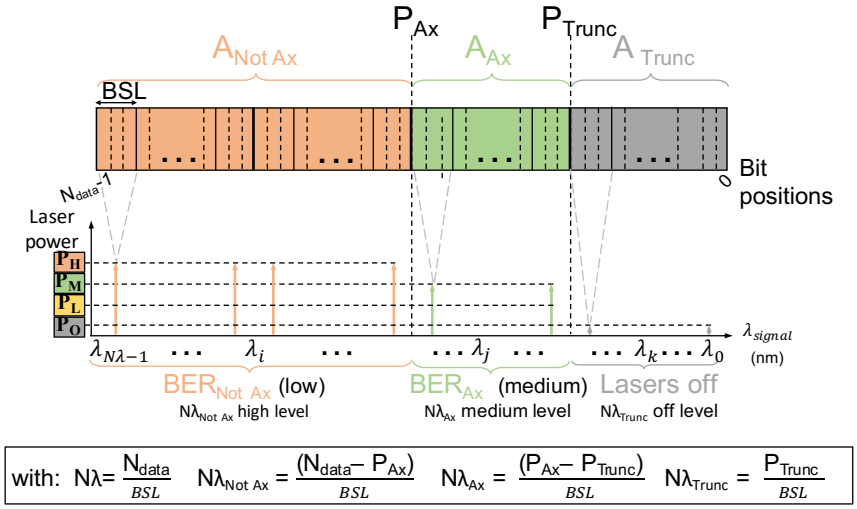


with: $N\lambda = \frac{N_{data}}{BSL}$ $N\lambda_{Not\ Ax} = \frac{(N_{data} - P_{Ax})}{BSL}$ $N\lambda_{Ax} = \frac{(P_{Ax} - P_{Trunc})}{BSL}$ $N\lambda_{Trunc} = \frac{P_{Trunc}}{BSL}$

Fig. 2. Example of data approximation scheme.

These 3 groups provide flexibility to adapts the robustness of the transmitted bits with respect to the bit significance. Table 1 summarizes the parameters used to identify and configure the different groups of bits of a data with the proposed method. While the flexibility can be controlled through the BER to reach and by selecting the bits areas to approximate or truncate (defined by $A_{Ax}$ and $A_{Trunc}$ on Fig. 2), in this paper we propose to fix these parameters based on the analysis of a set of applications. This analysis is presented in the next section.

| Parameters | Definition |
|---|---|
| $N_{data}$ | Data size |
| $P_{Ax}$ | Position of the higher approximated bit |
| $P_{Trunc}$ | Position of the higher truncated bit |
| $A_{NotAx}$ | Area of accurate bits |
| $A_{Ax}$ | Area of approximated bits |
| $A_{Trunc}$ | Area of truncated bits |
| $N\lambda$ | Number of wavelengths (nb of lasers) |
| $BSL$ | Bit Stream Length ($BSL = N_{data}/N\lambda$) |
| $BER_{NotAx}$ | BER for non approximated bits |
| $BER_{Ax}$ | BER for approximated bits |

Table 1. Summary of parameters used for the proposed method.

### 3.3 Power levels Definition

In this section, we present the method allowing to define laser power $P_H$, $P_M$, and $P_L$ according to the targeted BERs ($BER_{NotAx}$ and $BER_{Ax}$). As illustrated in Fig. 3, we plot three lines representing the received optical power with respect to the communication distance hop count for different injected optical power. The source optical power is estimated by:

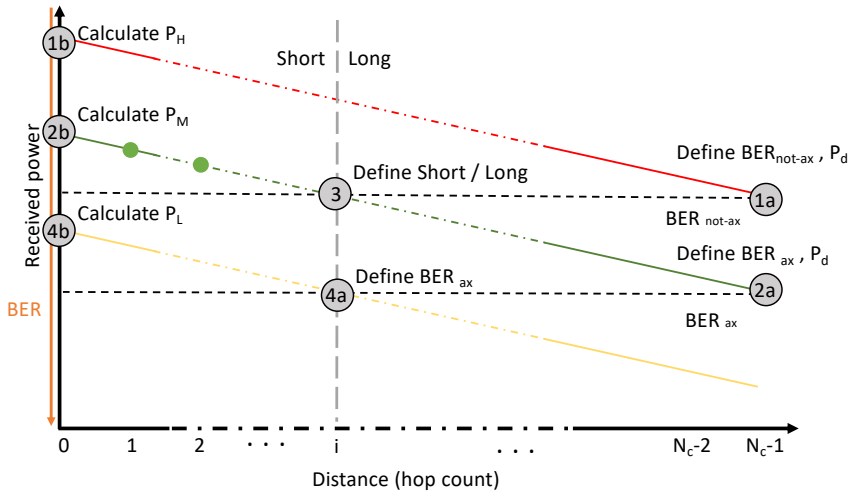


Fig. 3. *Short*/*Long* distance aware and power level computation method.

$$P_s \quad = \quad P_d + (Hop_{s,d} - 1) \times Loss_{mr} \times N\lambda + \; Hop_{s,d} \times Loss_h + \; Loss_d + Loss_{Xtalk} \tag{1}$$

Where $P_s$ is the optical power injected from source core (in dBm) and $P_d$ is the optical power received at the destinations core (dBm). $P_d$ depends on the targeted BER at the destination photodetector [37]. In this work, we defined the BER for approximate and accurate communications using system level simulation results, as detailed in section 4. The BER strongly depends on the attenuation experienced by the propagating optical signals, which depend on : i) the hop count $Hop_{s,d}$, ii) the number of wavelengths (or lasers) $N\lambda$ , iii) the MR through loss ($Loss_{mr}$, in dB), iv) the waveguide propagation loss per Hop ($Loss_h$, in dB/hop), v) the MR drop loss ($Loss_d$, in dB), and vi) the crosstalk loss ($Loss_{Xtalk}$, in dB). For this latter we consider the crosstalk model from [38]. The following details the method.

We first define the destination power $P_d$ allowing to reach the targeted $BER_{NotAx}$ at the furthest receiver (core at distance $N_c - 1$, mark (1a) in Fig. 3,). By taking into account the architecture size and the losses, we estimate the required optical power to be injected from the source $P_s$ using equation 1. The corresponding signal power is illustrated by the red line. Mark (1b) corresponds to $P_H$, the injected optical power to inject at the source core to guarantee error free communication for any source to destination in the architecture.

Similarly, we define the power level $P_d$ on the last receiver ($N_c - 1$) for the approximate communication obtained using $BER_{Ax}$, mark (2a). From equation 1, the required optical power at the source allows defining $P_M$, mark (2b). As represented by the green line, the optical signal power for approximate communication remains high enough to reach an accurate communication for short range communication. From this observation, we define a limit between short range and long range communication, mark (3), which correspond to the highest hop count from which $P_M$ allows to reach $BER_{NotAx}$. This will allow to reduce the laser power for short range communication while maintaining robust data transmission. Finally, in order to efficiently transmit approximate data, we define $P_d$ (mark (4b)) as the optical power allowing to reach $BER_{Ax}$ for all short range communication, mark (4a).

The method allows us to define the optical power to be emitted by the lasers and to classify each communication as short or long distance. The power levels and hop threshold are then used in the ONIs as parameters to configure the laser drivers regarding the communication distance and the data to transmit.

### 3.4 Power saving

Fig. 4 illustrates six possible communication scenarios to save energy for the transmission of 32 bits data. For this scenario, 8 wavelengths ($\lambda_7$ to $\lambda_0$) are used to simultaneously transmit 8 groups of 4 serialized bits. Each scenario is noted as xNA/yA/zT with: x Not Approximated, y Approximated, z Truncated bits. Regarding long distance communication (Fig. 4-a), we first illustrate a scenario where only approximation is used, assuming $A_{NotAx} = 12$ and $A_{Ax} = 20$. This scenario is noted as $12NA/20A/0T$. This leads to lasers at $\lambda_7$ to $\lambda_5$ and $\lambda_4$ to $\lambda_0$ configured with $P_H$ and $P_M$ respectively. Second scenario, $12NA/0A/20T$, involves only truncation ([13]) which refers to $A_{NotAx} = 12$ and $A_{Trunc} = 20$ and it leads to laser configuration $P_H$ for $\lambda_7$ to $\lambda_5$ and $P_{Off}$ for lasers at $\lambda_4$ to $\lambda_0$. Third scenario, $12NA/8A/12T$, is a combination of approximation and truncation with $A_{NotAx} = 12$, $A_{Ax} = 8$, and $A_{Trunc} = 12$ which leads to $P_H$, $P_M$ and $P_{Off}$ for lasers at $\lambda_7$ to $\lambda_5$, $\lambda_4$ to $\lambda_3$ and $\lambda_2$ to $\lambda_0$ respectively. Moreover, when all of these 3 scenarios are applied in short-range communication, the laser power levels are reduced accordingly ($P_H \rightarrow P_M$ and $P_M \rightarrow P_L$) (as presented in the three scenarios of Fig. 4-b).

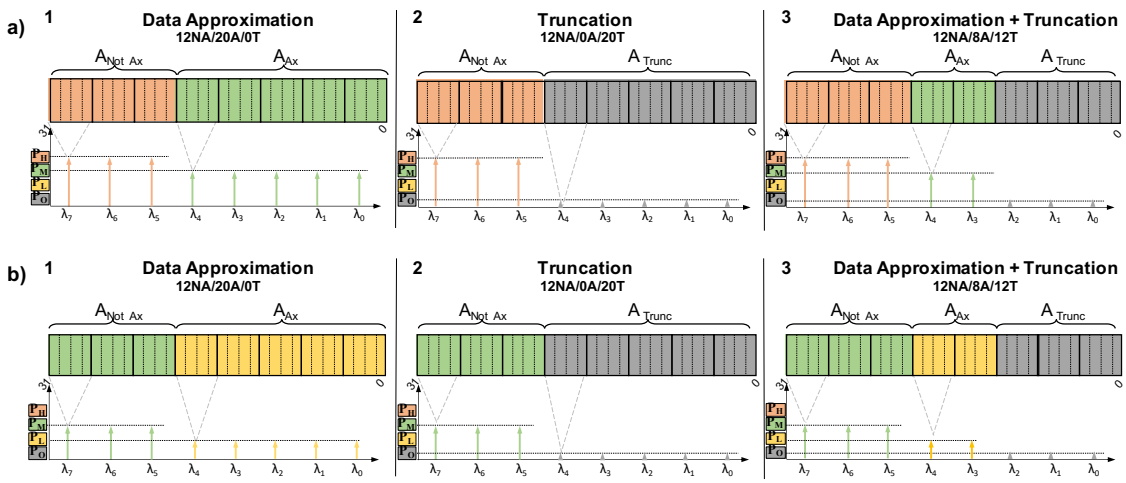


Fig. 4. Laser power level configurations to transmit approximated floating-point data assuming a) Long distance and b) Short distance communications. $N_{data} = 32$, $N\lambda = 8$ and $BSL = 4$. xNA/yA/zT with x Not Approximated, y Approximated, z Truncated bits

Compared to the scenario where all bits are always transmitted at full power, regardless of the communication distance, our approach allows saving power, as detailed in Equation 2. The parameter values used in equation 2 are computed from equation 1 that considers the required power at the receiver and the propagation losses. To calculate power gain, power consumption from proposals $P$ over Baseline scenario $\lambda \times P_H$ has to be computed. For baseline laser power is configured as $P_H$ for all the wavelengths to guarantee accurate communication for all the distances, equation 2a. Moreover, in equation 2b, to obtain power consumption from proposal, power consumption on three bit areas (defined in section 3.2) have to be computed. First, MSBs of FP $[N_{data} - P_{Ax}]$ is configured with $P_H$ or $P_M$ level according to the Short and Long distance. Secondly, approximate bit area $[P_{Ax} - P_{Trunc}]$ is configured with $P_M$ or $P_L$ to

target approximate communication. Last, for the lasers that are in charge of Truncation bit area $[P_{Trunc} - 0]$ are set in off state $P_{Off}$. The summation of all the power consumption from those bit areas defined the total power consumption for the proposal. The actual power saving depends of bit area configuration, targeted BER and distance configuration.

$$PowerGain = \frac{P}{N\lambda \times P_H} \tag{2a}$$

$$\tag{2b}$$

$$\begin{aligned} \text{with } P &= \frac{1}{BSL}\Big[(N_{data} - P_{Ax}) \times P_1 \\ &\quad +(P_{Ax} - P_{Trunc}) \times P_2 \\ &\quad +P_{Trunc} \times P_{Off}\Big] \\ \text{and with } (P_1, P_2) &= \begin{cases} (P_H, P_M) \text{ for long distance,} \\ (P_M, P_L) \text{ for short distance.} \end{cases} \end{aligned}$$

## 4 EXPERIMENTAL SETUP

In this section, we present the system parameters for the architecture and simulation platform that are considered for the evaluations performed in the next section.

### 4.1 Considered 3D architecture

Table 2. Hardware parameters of the simulated architecture

| **Parameters** | **Values** |
|---|---|
| Nb of clusters | 16 $(4 * 4)$ |
| Nb cores by cluster | 4 |
| ONI distance | 10 $mm$ |
| L1 I / D caches | 64 $KB$ each private |
| L2 cache | 512 $KB$ private |
| L3 cache | 4 $M$ shared (in a cluster) |
| Cache protocol | MSI |

For experimental setup, we consider an architecture based on two layers: one electrical layer to support the computation part, and one optical layer for the optical communications (as presented on Fig. 5-a). The chip size considered for this section is equal to $1600mm^2$, hence the distance between two consecutive ONIs is therefore equal to 10 $mm$. As summarized in Table 2, the electrical layer is composed of 16 clusters of 2 by 2 cores. These cores include a last level cache (L3) shared with the others. Each core has its own private L1 data (L1d), instructions (L1i) caches and L2 cache. Furthermore, a MSI protocol ensures coherency between the distributed caches. All these memory organisation parameters are listed in Table 2. Every clusters are connected with ONIs through TSVs. The ONI blocks are located in optical layer and have 1 transmitting and 15 receiving waveguides which are featured with SWMR and WDM (as depicted in Fig. 1). Each waveguide transmits 8 optical signals on different wavelengths, hence each having a 4 bits stream data (i.e. $BSL = 4$ and $N\lambda = 8$).

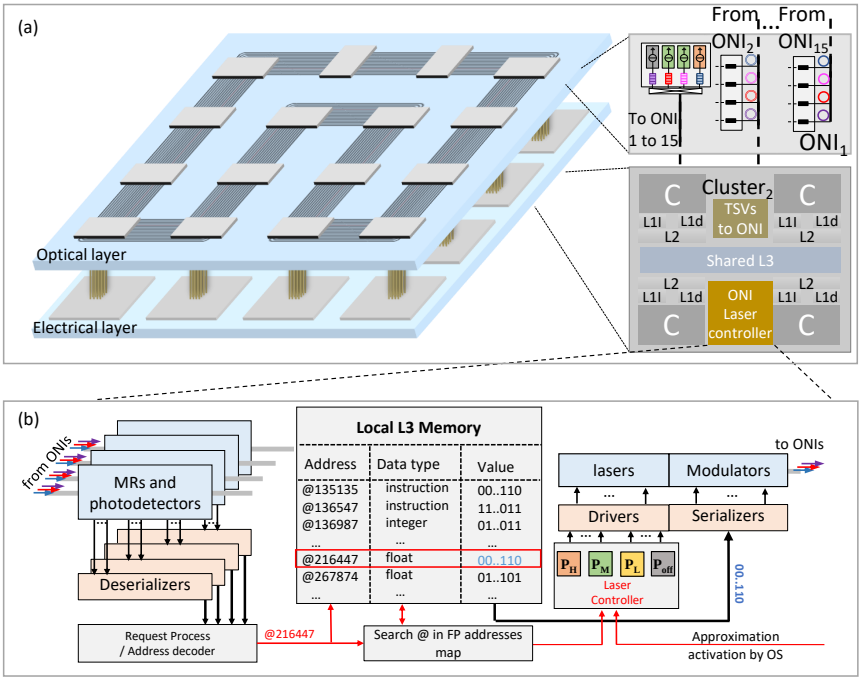


Fig. 5. a) Hardware Architecture (16 clusters of 4 cores) b) SWMR channel management for approximate communication

### 4.2 Approximation under Cache Coherence Traffic

Cache coherency involves traffic between cache memory and last level of cache that is mainly initiated by cache misses or write-back accesses [39]. However, the traffic initiated by cache coherence protocols also involves addresses of the data to be transmitted and is thus data type agnostic. We assume that approximation are enabled by the OS based on user requirements. This is achieved by setting or resetting flag "Approximation activation by OS" at run-time (figure 5). In case approximations are disabled, all communications are carried without errors using $P_H$ and $P_M$ for long-range and short-range communications respectively. Otherwise, the lasers are configured as follows. First, we assume that i) accurate data and approximate data are located in different address spaces of the L3 memory and ii) ranges of addresses associated to a given data type are stored in local tables. Then, from the packet destination ID, the ONI evaluates whether the communication distance is short or long (see details in section 6.1). Both information (accurate/approximate and short/long) allow the ONI manager to configure the laser power level from the laser configuration table. The details of communication between OS and ONIs is out of scope for this paper and it is studied in [40]. By considering all these parameters, Laser controller configures laser powers at the different levels ($P_H$, $P_M$, $P_L$ and, $P_{Off}$). Thus, cores have two ways to communicate, one without error and one with approximation mode. The Fig. 5-b illustrates the principle with a cache miss scenario. The L3 on cluster 2 receives a request containing the address of a data to be returned. The address (@216447 in the example) is compared to the address ranges of data to be approximated (FP numbers). In case of matching, the power levels of lasers corresponding to the approximated and truncated bits are configured. To limit the hardware complexity, the proposed architecture only integrates one approximation and truncation configuration (see conclusion of section 6.3 for more details) and can be activated at run-time by an Operating System with a flag signal

[40, 41]. The manager then grants the SWMR access to the L3 that initiated the request and the data is transmitted to the serializers for signal modulation. In order to reduce the latency and hardware overhead, the whole cache line is transferred: the size of response packets (which can be approximated) is larger than packet request size (which can never be approximated) as it contains data. We assume the use of a packet header integrating the IDs of source and destination cores, data address, data type, and flag indicating if the transmitted data are approximated. The header is always transmitted with high laser power, i.e. without approximation.

### 4.3 Simulation platform

For evaluation, we consider error tolerant applications from Approxbench [11]. Table 3 presents the four benchmark applications used in this paper: Blackscholes, Fluidanimate, Canneal and Streamcluster performing financial algorithms, animation, routing and online clustering, respectively. Moreover, Approxbench provides the output error metric for each application, which allows to estimate the Quality of Result (QoR).

This output error is the difference on error metric between the application with and without approximate communications. As each application performed has specific goal, each one has a particular error metric which is defined in Table 3. For instance, Blackscholes is financial pricing application that predicts price of the market based on previous market values and trend. Thus, the output error metric is the Mean Square Error (MSE), and is computed on mean error of output values as follow in Equation 3:

$$MSE = \sqrt{(\frac{1}{n})\sum_{i=1}^{n}(Accurate_i - Approximate_i)^2} \tag{3}$$

Based on this equation, approximated output values $Approximate_i$ are compared with original ones $Accurate_i$ to determine output error.

These applications are simulated in multi-core simulator Sniper [16]. From the simulations, we extracted the output results in order to compute the output errors and also the communication traces linked to the traffic between clusters in NoC. From the communication traces, which embed information of packet type, source, destination and sent time, we analyze communication distribution and type of data.

Table 3. Used Approximate Benchmark and associated error metric [11]

| Applications | Descriptions | Error metrics |
|---|---|---|
| Blackscholes | Financial pricing | Mean error |
| Canneal | VLSI routing | Routing Cost |
| Fluidanimate | n-body simulation | Distance in elements |
| Streamcluster | Online clustering | Cluster center distance |

### 4.4 Benchmarks and traffic analysis

To evaluate the impact of the proposed method, we analyze the traffic distribution for the selected applications. This analysis includes both hop distance and data type, within the considered 16 clusters architecture.

Fig. 6-a) presents ratio of communication by hop count. We can note that the communication traffics over the network are well balanced for all the hop distances and for any application. From this well distributed communication, we can

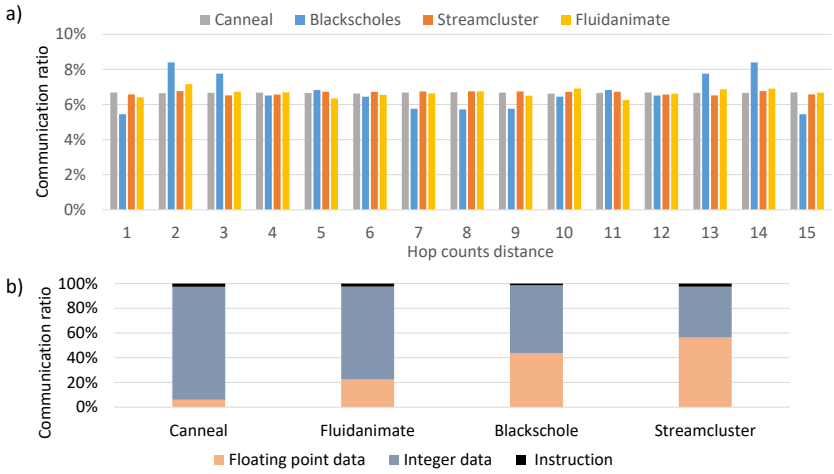


Fig. 6. Traffic distributions on 16 clusters architecture: a) by hop count b) by data type.

say the packets are uniformly exchanged among the cores. Secondly, Fig. 6-b) presents the traffic distribution by data type. The communications are classified in Instruction, Integer, and Floating-Point data. The ratio of type of data is linked to the application and it is defined at the application designing level. Only Floating-Point (FP) ratio refers to potential data type that may benefit from power saving by using the proposed approximate method. We can see that Canneal only have around 4% of FP data communication, while the others reached 22%, 44% and 60%, for Fluidanimate, Blacksholes and Streamcluster, respectively. From this analysis, it is evident that applying energy reduction technique on FP data for Canneal will not produce a considerable energy reduction, unlike the other applications, and especially Streamcluster.

## 5 IMPACT OF APPROXIMATION AND TRUNCATION ON QUALITY OF RESULT

In this section, we first investigate the impact of bit truncation on the Quality of Result (QoR). Then we analyze the QoR results when bit approximation and bit truncation are combined together.

### 5.1 Truncation impact on QoR

The truncation is the most straightforward technique to save energy, as certain amount of bits are not sent. However, it may also impact most the output of application, hence degrade the QoR. Fig. 7 illustrates the output error effects when only bit truncation is applied on FP data, starting from the LSB bit. It has to be noticed that we increased the number of truncated bits four by four up to 12 bits, as below that value no effect has been sensed, then we increased this number one by one up to 24 bits, that exhibits too much error for every application. We have set an output error threshold lined at 10% which is widely used as an acceptable error range for approximate computing [11, 14, 15, 30]. From the analysis, it is certain that applications differently react to a given approximation level. As an example, Canneal does not generate any visible output error with truncation. It may be explained as Canneal contains only 4% of FP data usage (Fig. 6-b). Approximation on very small quantity of FP data does not critically impact the whole application. Blackscholes and Streamcluster show low output error similar to Canneal until truncation of 23 bits. However, at 24 bits truncation, the produced output errors exceed the 10% threshold. Fluidanimate, meanwhile, produces a large output

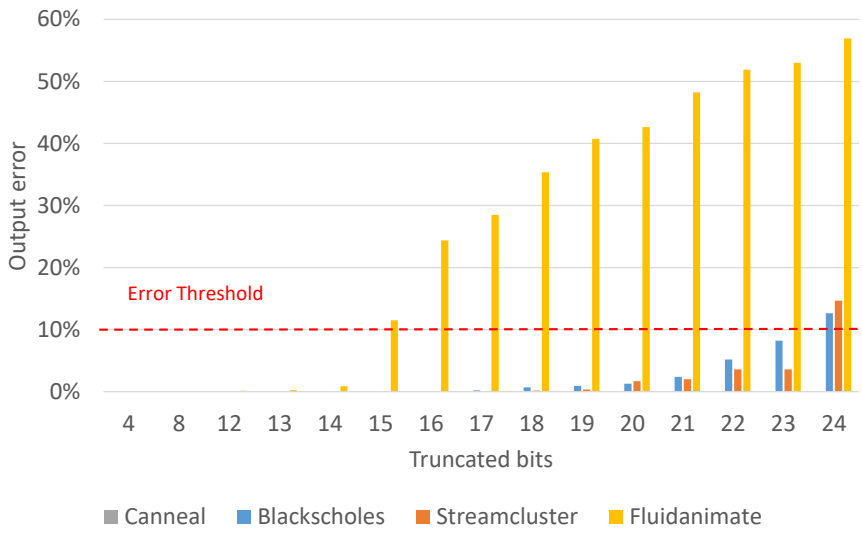


Fig. 7. QoR according to the number of truncated bits for 32 bits floating-point data.

error, quickly reaching the 10% threshold, starting from 15 truncated bits. This application contains 21% of FP data which is quite low compared to Blackscholes and Streamcluster. However, unlike Canneal, Fluidanimate only tolerates a low approximation level due to data sensitivity. For Blackscholes and Streamcluster, even if they contain 44% and 58% FP data, they still remain under acceptable error threshold as FP data in those applications are not sensitive and have less impact on results.

We can notice that all the applications suffer the 24 bits of truncation. Indeed, in FP single precision presentation, LSB bits from 0 to $23^{th}$ are mantissa, and bits from $24^{th}$ to $31^{th}$ are exponent which has significant meaning for the data after the number normalization. Therefore, truncating $24^{th}$ for all the applications leads to large amount of error. For these experiments, we consider an ONI with 8 lasers serializing a group of four bits ($BSL = 4$). Thus, it is possible to apply truncation by steps of 4. For instance, truncation of 20 bits or 24 bits is configurable in this architecture. Therefore, in our system, the most beneficial configuration just with truncation, under the 10% threshold, is 20 bits for $A_{Trunc}$ and 12 bits for $A_{NotAx}$.

### 5.2 Impacts of truncation and approximation combination

In the previous subsection 5.1, we found that truncating bits close to MSBs produces large amount of errors, in particular when the truncation is done on bits coding the FP exponent. In order to push further the energy reduction, approximation instead of truncation on bits close to MSB should be less aggressive for results degradation. Indeed, approximation of bit is less aggressive than truncation from application point of view. It is possible by adapting laser power level according to targeting BER and produces power saving.

The Fig. 8 present how combinations of both truncation and approximation techniques impact the QoR of applications. The output errors are presented with different bit range parameters: Not Approximated bits (NA), Approximated bits (A) and Truncated bits (T). The explored combinations range from 8 MSBs not approximated (sign and exponent of 32 bits FP data are not approximated) and the remaining bits approximated (noted as 8NA/24A/0T), to 8 MSBs not approximated and only truncation (noted as 8NA/0A/24T). The intermediate values are increased by steps of 4 more

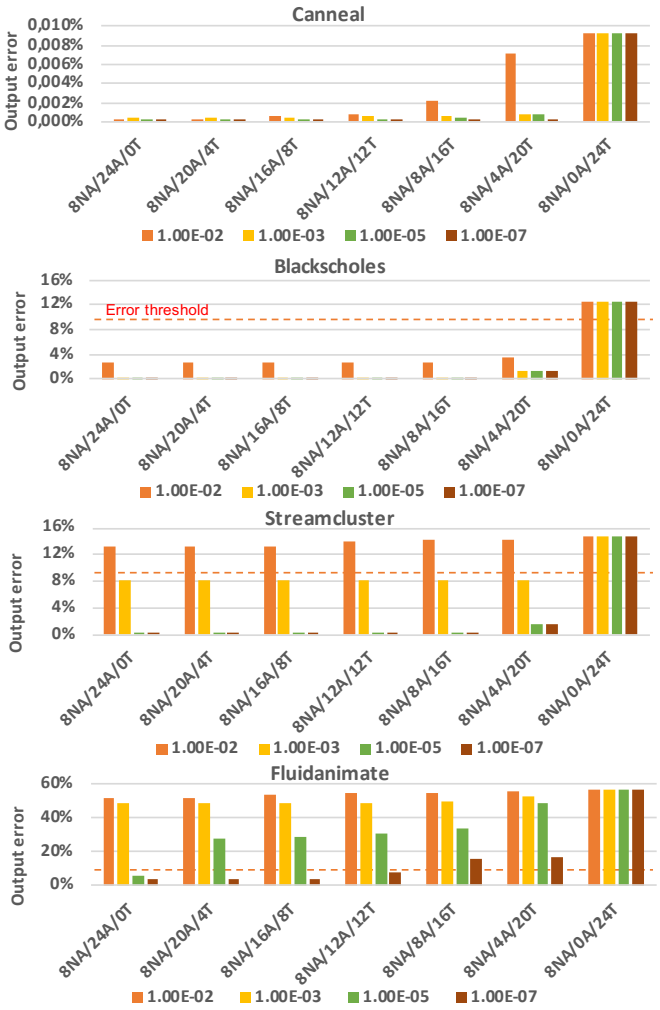


Fig. 8. Combination of Truncation and Approximated bits considering output error. xNA/yA/zT (with x Not Approximated, y Approximated, z Truncated bits).

truncated bits instead of 4 approximated bits. The step 4 is linked to the bit stream length ($BSL = 4$ in our experiments). Regarding the approximated bit levels, we explored the aforementioned combinations with different $BER_{Ax}$ values: $10^{-2}$, $10^{-3}$, $10^{-5}$, and $10^{-7}$.

Firstly, from the results shown in Fig. 8, we can clearly see that the combination of mixing truncation and approximations provides a better QoR, for any BER value, than only 24 bits of truncation (represented by 8NA/0A/24T). As already explained in the previous section, applying 24 bits truncation impacts the first bit of the exponent within the Floating-Point number. Secondly, we can see that increasing BER value decreases the QoR, due to more error on bits. For instance, if we consider Fluidanimate, for 24 approximated bits (i.e. 8NA/24A/0T), the QoR goes from 2.5% to 5%, and from 48% to 50% of output error, respectively with a BER increasing from $10^{-7}$ to $10^{-5}$, and from $10^{-3}$ to $10^{-2}$. This

is also observable for any combination of approximation with truncation. Thirdly, we can notice that, for any BER value, the QoR is more degraded as the number of truncated bits increased. This is especially visible for Fluidanimate which is more sensitive to errors: for a BER value of $10^{-5}$, the output error goes from 5%, to 58%, while passing by 30%, respectively with only 24 approximated bits (8NA/24A/0T), only 24 truncated bits (8NA/0A/24T), and half approximated and half truncated (8NA/12A/12T).
Regarding the best combinations to keep an acceptable output result:

- Canneal can go for 24 truncated bits. As it is not a Floating-Point sensitive application, the combination of techniques does not provide any advantage.
- Blackscholes exhibits less than 10% of errors for any combination except for 24 truncated bits. Hence, it can even apply strongly degraded BER of $10^{-2}$. As it is evident that truncation of 4 bits consumes less power than 4 approximated bits, we may expect that a combination with 4 approximated bits and 20 truncated bits is the best solution if we target the less consuming solution with an acceptable QoR (8NA/4A/20T with a BER of $10^{-2}$) for this application. A complete analysis on QoR and power savings is presented on the next section.
- Streamcluster has somehow the same behavior as Blackscholes, however the BER should be better than $10^{-2}$. Hence, the best combination should be for 8NA/4A/20T with a BER of $10^{-3}$
- Fluidanimate, is more sensitive to data errors. For any combination, if the BER is worse than $10^{-5}$, the output error is more than 40%. The combination with the higher number of truncated bits providing less than 10% of output error is for 12 approximated bits at a BER of $10^{-7}$ along with 12 truncated bits (8NA/12A/12T).

From these experiments, we can conclude that every application reacts differently to the approximation and truncation. However, we can see three trends. First, some applications are not really sensitive, like Canneal. However, this benchmark only has less than 5% of Floating-Point data in its communications. Hence, the opportunity for energy saving is not relevant. Secondly, some benchmarks are very sensitive to approximate, like Fluidanimate. Hence, this kind of application requires a specific and accurate combination of approximation and truncation. Finally, it is possible to find some applications with a non negligible amount of Floating-Point data that can be approximated using the same combination of parameters. This latter is the case for Blackscholes and Streamcluster.

From this analysis, the proposed method can be used for two purposes:

- Application specific communication interconnect: adapting i) the number of approximated and truncated bits and ii) the BER for every application. However, this will require a complex ONI that will have to adapt on run time to the application needs;
- Generic interconnect for approximation: it is possible to determine one combination of truncated and approximated bits that suits the most of approximable applications. However, by going in that way, the ONI will not be able to approximate some applications that then will be run using non approximated communication.

For the next section, we will explore the energy gains provided by approximation and truncation while also considering the QoR. We consider for the following a generic interconnect that use the following scheme for approximating floating-point data: 20 truncated bits with 4 approximated bits (8NA/4A/20T) sent with a $BER_{ax}$ of $10^{-3}$. This solution will then be compatible with Blacksholes and Streamcluster applications, while not possible to be applied on Canneal and Fluidanimate.

## 6 POWER CONSUMPTION VS QOR

This section analyzes the power saving according to i) the level of approximation and truncation, ii) the targeted BER for approximate and accurate communication, and iii) the communication traffic between processors and memories. For this purpose, we first use the proposed method to define the laser power levels and the limit between short-range and long-range communications, which are key parameters in our system. For this purpose, we explore the impact of the distance between the ONI (system level parameter) according to the waveguide loss (technological level). Then, we provide a detailed analysis of the power contribution for each type of communication. Exploration of QoR - power saving tradeoff is provided and we compare the found solutions with the ideal ones. As discussed in the previous section, not all the applications from Approxbench lead to significant power saving under approximate communications schemes. Nevertheless, for the sake of reduced design complexity, we decided to use the same approximation scheme, i.e. 8NA/4A/20T configuration with $BER_{Ax} = 10^{-3}$, for all the applications. In the following, we mainly focus our studies on the StreamCluster application. Finally, we propose a comprehensive study on the design complexity of the drivers for our method and Lorax [13].

### 6.1 ONI Specifications

In the following, we design the ONIs according to system level parameters defined in Table 2 and by considering $BER_{NotAx} = 10^{-12}$ and $BER_{Ax} = 10^{-3}$ (defined from previous simulation results). Our goal is twofold: first we aim to define power levels $P_H$, $P_M$ and $P_L$. Second, we aim to define the hop count from which long-range communications start. We first assume minimum optical power of -8 $dBm$ and -12 $dBm$ for accurate and approximate data reception respectively [37]. We assume 0.02 $dB$ MR through loss, 0.7 $dB$ MR drop loss, [42] an FSR of 8$nm$ and a MR quality factor of 20,000 [43]. We consider distance between ONIs ranging from 0.8 $cm$ to 1.3 $cm$ and consider waveguide losses of 1 $dB/cm$ [44] and 0.25 $dB/cm$ [45]. While 0.25 $dB/cm$ corresponds to ridge waveguides not compatible for side-coupling with microrings [46], our aim is to demonstrate the ability of the method to also consider prospective propagation loss and to carry our comparison with Lorax. Fig. 9 a) reports the exponential growth of the laser powers for increasing distance between ONI, assuming the two considered propagation losses. $P_H$, $P_M$ and $P_L$ are estimated based on the method defined in Section 3.3. The losses breakdown to communicate with the last receiver are shown in Fig.9 b) for a number of wavelengths ranging from 8 to 24 and for a 0.25dB/cm waveguide loss. The waveguide loss, which dominates the total losses for a small number of wavelengths, is independent from the number of lasers since it only depends on the communication distance. On the contrary, the MR through losses linearly increase with the number of wavelengths due to the higher number of crossed MRs. Regarding the crosstalk losses, it only represents 1.3% for 8 wavelengths. Indeed, for an FSR of 8$nm$, the channel spacing is 1$nm$ (~120GHz), which results in negligible crosstalk losses as demonstrated in [43]. For a same FSR, increasing the number of wavelengths to 24 leads to 7.9% of losses provided by the crosstalk, which is acceptable and can be lowered by increasing the channel spacing [38]. In the rest of the study, we will consider 8 wavelengths. $P_H$, $P_M$ and $P_L$ calculations are detailed for a distance between ONIs of 1 $cm$ in Fig. 9 c) (0.25 $dB/cm$ propagation loss) and Fig. 9 d) (1 $dB/cm$ propagation loss). As seen in Fig. 9 c) (respectively Fig.9 d), to reach $BER_{NotAx}$ at ONI 15, the higher transmission losses are compensated by an increase of $P_H$ up to 707$\mu W$ for 0.25$dB/cm$ (respectively up to 9549$\mu W$ for 1 $dB/cm$). Similarly, to reach $BER_{Ax}$ at ONI 15, $P_M$ of 281$\mu W$ is needed for 0.25 $dB/cm$, which leads to $BER_{NotAx}$ at ONI 5, while $P_M$ must be fixed at 3801$\mu W$ for 1$dB/cm$, which leads to $BER_{NotAx}$ at ONI 11. Hence, short range communications for 1 $dB/cm$ include ONIs 1 to 11, against ONIs 1 to 5 for 0.25 $dB/cm$. In other words, the total losses experienced by the optical signal for hops 11 to 15 under 1 $dB/cm$

are similar to the losses for hops 5 to 15 under 0.25 $dB/cm$. For comparison purpose with [13], we consider power values obtained for 0.25 $dB/cm$ in the rest of the paper ($P_H = 707\mu W$, $P_M = 281\mu W$ and $P_L = 112\mu W$), as summarized in Table 4.

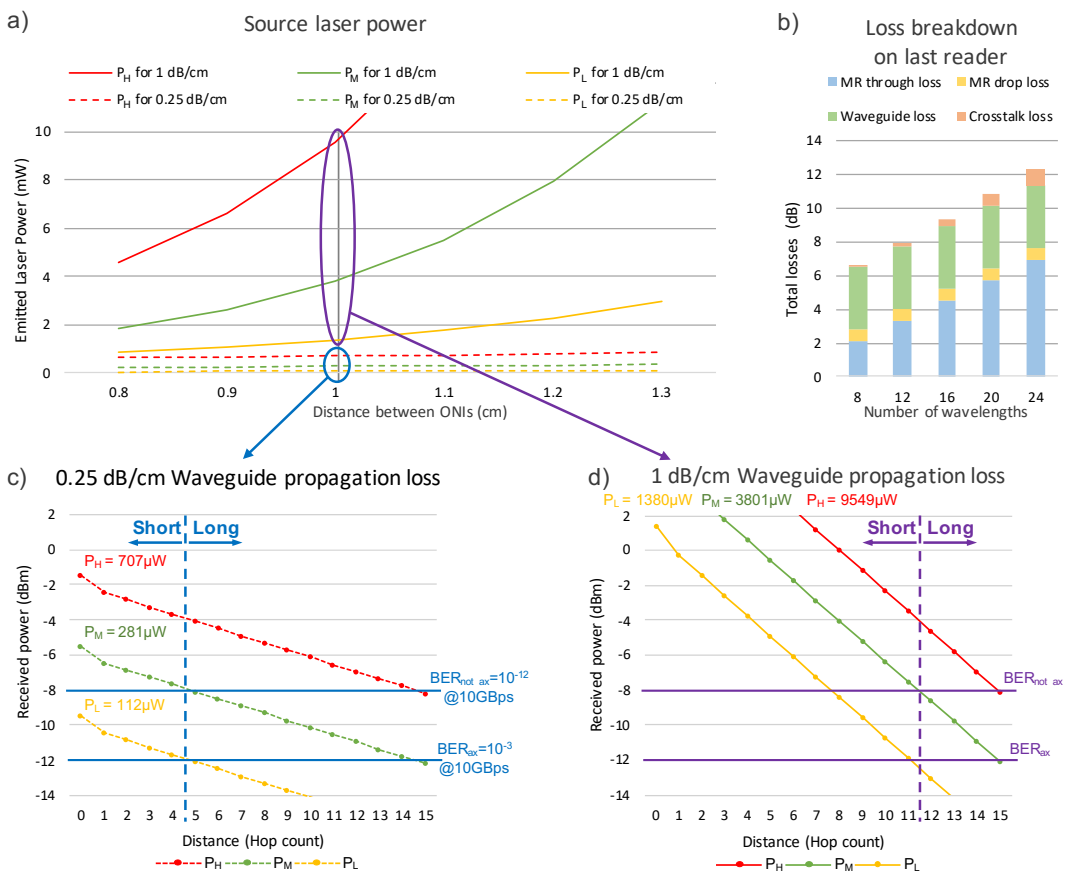


Fig. 9. a) Required laser output power according to the distance between ONIs and for $BER_{ax} = 10^{-3}$ and $BER_{notax} = 10^{-12}$. b) Communication losses overview to reach the last receiver with 0.25dB/cm waveguide propagation loss for 8 to 24 wavelengths. Details of the method to define short/long-range communications and to calculate $P_H$, $P_M$ and $P_L$ for a ONI distance of 10$mm$ are given for c) 0.25dB/cm and d) 1dB/cm waveguide propagation losses.

### 6.2 Power saving analysis

Fig. 10 shows the power saving for existing methods [12, 13] and the proposed one. All the results are compared with the optical power of the *Baseline* configuration, which only use the highest optical power level ($P_H = 707\mu W$), without regards on data type and distance of communication. The floating-point data are represented with orange colors, while the rest of communication (integer data and instructions) are represented with blue colors. The light orange (respectively light blue) represents the contribution of short distance communications (respectively long distance

Table 4. ONI design parameters

| Parameters | Descriptions | Typical values |
|---|---|---|
| | **Technological parameters** | |
| $Loss_{mr}$ | MR through loss | $0.02\,dB$ [47] |
| $Loss_d$ | MR drop loss | $0.7\,dB$ [47] |
| $Q$ | MR Quality Factor | $20,000$ [43] |
| $Loss_h$ | Waveguide loss per hop | $0.25\,dB/cm$ [42] |
| $P_d$ | Optical power received by photodetector | $-8\,dBm$ at $10^{-12}$ BER<br>$-12\,dBm$ at $10^{-3}$ BER |
| $FSR$ | Full Scale Range | $8nm$ [43] |
| $Laser_eff$ | Laser efficiency | 0.33 [48] |
| $B_R$ | Bit-rate | 10Gb/s |
| | **ONI Parameters** | |
| $BER_{NotAx}$ | BER without approximation | $10^{-12}$ |
| $BER_{Ax}$ | BER with approximation | $10^{-3}$ |
| $P_H$ | High power | $707\,\mu W$ |
| $P_M$ | Medium power | $281\,\mu W$ |
| $P_L$ | Low power | $112\,\mu W$ |
| $N\lambda$ | Number of wavelengths | 8 |
| Short distance | Nb of hops from | 1 to 5 hops |
| Long distance | source to destination | 6 to 15 hops |

communications). As previously defined, the threshold between long and short distance communications is equal to 5. Considering this distance threshold, the *Baseline* shows that 33% of transfers are short communications, while 66% are long communications. Regarding the data type, up to 58% can be transmitted using approximation (see FP ratio in Fig. 6-b). The second stacked vertical bars, named $Short/Long$ on the figure, show the power contribution when the lasers power levels are adapted according to Short/Long distances. In this case, the long distance communications are sent with the highest power $P_H$, while the power is reduced to $P_M$ for short distance communications and for any data type. This results in 20% of power saving compared to the *Baseline*. The last method that only benefits from distance consideration for power saving is *Distance proportional* shown with the third stacked bars. It assumes to only use the required power for any distance within the architecture [13], hence it requires as many power levels as distance hops available. We can observe that the power required for both short and long distances of communication are reduced. It allows to save 21% and 41% more power compared to $Short/Long$ and *Baseline* methods, respectively. It shows most effective power efficiency compare to *Baseline* and $Short/Long$ as it avoids to waste any extra power for the transmission. However, the implementation of fully distance proportional links requires significant design efforts for the laser CMOS drivers and leads to hardware overhead due to the number of laser power level to be pre-configured. Moreover, such links do not scale well as the hardware complexity linearly increases with the number of destinations and configurable BER. Indeed, assuming N cores architectures and M BER levels, the approach leads to $(N-1) \times M$ different power levels. On the contrary, the approach we propose only requires 3 power levels. Overall, a comprehensive power analysis is needed to further compare the approaches by considering the overheads of the optoelectronics and electronics components. The detail of this laser design complexity for fine grain multi power level are studied in section 6.4.

Scenario $12NA/20A/0T$ involves approximation for 20 LSB of FP data, with a BER of $10^{-3}$, hence by setting laser power to $P_M = 281\,\mu W$ and without distance management [12]. The system benefits 21.8% power saving compared

to *Baseline*, which is due to a 62% reduction in the power needed to transmit FP data. Obviously, no power saving is obtained for non-FP data. By assuming $Short/Long$ distance management for the previous method, approximated bits will be sent using $P_L = 112\,\mu W$ and $P_M = 281\,\mu W$ for short and long distances, respectively. An overall power reduction from 78.1% to 63.6% is achieved. Power consumption for communications that only use truncation to save power, without distance management, is presented in $12NA/0A/20T$. In this result, 20 LSB are truncated on FP data communication by turning off lasers. The architecture benefits from 36% power saving compared to *Baseline* throughout all communications. Moreover, if $Short/Long$ distance management is associated to this solution, hence leading to $12NA/0A/20T + Short/Long$ result, it leads to an extra 13% power saving, then reaching a total power consumption of only 50.9% versus the *Baseline*.

By combining truncation and approximation without considering the distance, scenario $8NA/4A/20T$ shows 40% of power reduction. Associated to $Short/Long$ distance management, ($8NA/4A/20T + Short/Long$), the power required for the communication drops to 47%, thus leading to a significant improvement. All these results demonstrate the needs to combine approximation techniques at the levels of both optical link and data representation.

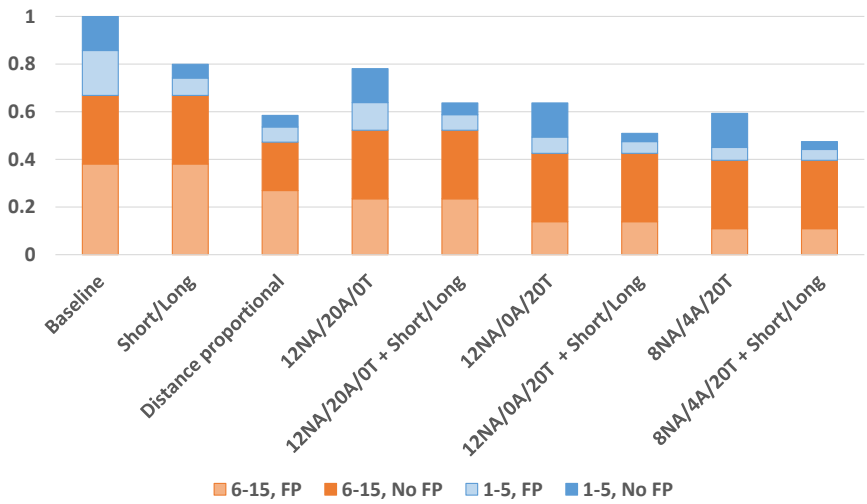


Fig. 10. Normalized power consumption for the considered communication schemes.

### 6.3 Power Vs Quality of Results Trade-offs

To fully evaluate the different techniques, we have to consider both power consumption and application output error. The Fig. 11 highlights the design space with all the possible techniques to handle approximate communications in ONoC (including the following different $BER_{Ax}$: $10^{-2}$, $10^{-3}$, $10^{-5}$ and $10^{-7}$). The bottom of the figure shows zooms of valid solutions from the top figure. All the optical power consumptions have been normalized to the *Baseline* represented by the red cross. This baseline only uses $P_H$ power level, hence no distance management and no approximation are considered. For more clarity, we grouped the results by different communication distance managements: *No Distance aware*, $Short/Long$ and *Distance proportional*, in red, blue, and green, respectively.

In each distance management group, the solutions that do not consider approximation or truncation are represented with a cross. Moreover, in each group, some points are located on Pareto front lines which indicates the most effective

results in terms of power consumption and output error. The solutions on Pareto front lines are represented with triangles and the configuration are highlighted, e.g. 4A/20T $10^{-3}$ stands for 4 approximated bits, 20 truncated bits, and $BER_{ax} = 10^{-3}$. For more clarity in the Fig. 11, the number of non-approximated bits is not shown, but can be easily deduct as the total number of bits of FP data are 32bits. We also represented the solutions that provide more than 10% of output error, there are located on the right gray part.

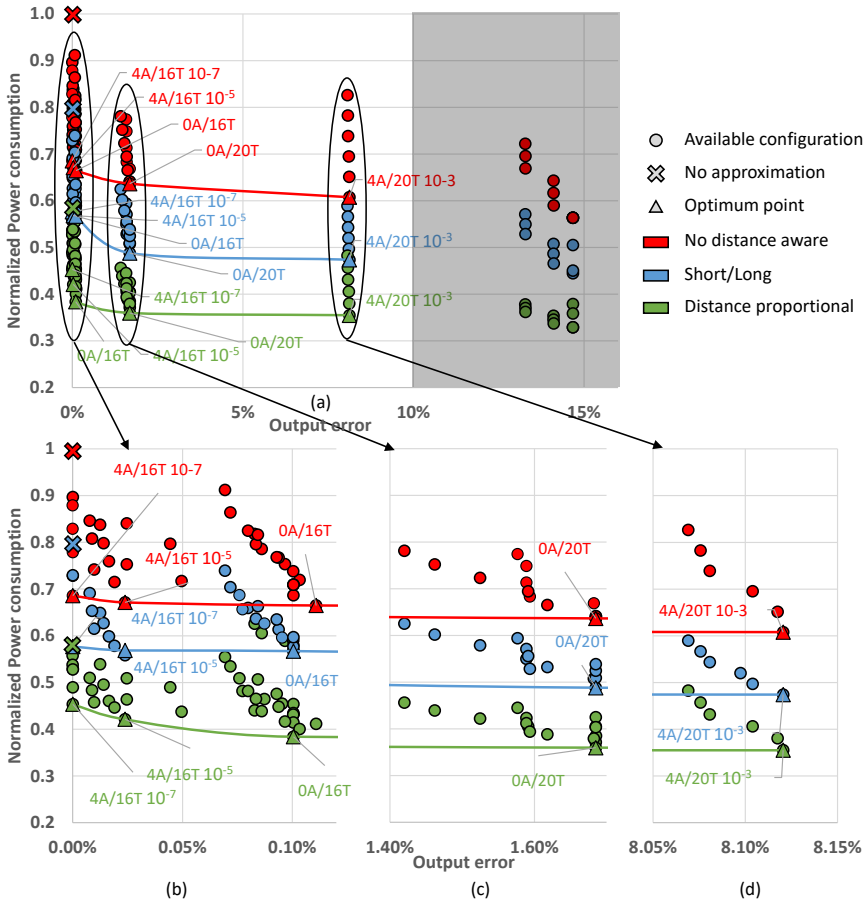


Fig. 11. Power consumption and output error trade-off: (a) design space, (b), (c), and (d) are zooms.

First of all, we can see that there is a huge exploration space with high differences in terms of power versus QoR. The way to manage the distance of communication provides a shift from 12% for the same configuration in terms of approximating or truncating data. Unsurprisingly, the distance proportional solutions are the less power consuming, but at the cost of high complexity. It is interesting to note that, for distance proportional solutions, some configurations can consume more that the best solution with only two distances considered. The same trend can be seen for two distances management that has solutions worse than the Pareto front of the *No Distance aware*.

Regarding the best solutions in terms of trade-offs between power consumption and output errors, we see the same trends on the different color groups. The first main cloud of solution, zoomed in Fig 11-b, are dominated with three

points on the Pareto fronts in each group. All these solutions provide less than 0.15% of output errors. The solutions $16NA/0/16T$ with a maximum of 16 truncated bits provide a very small degradation of the output error, but allow to save a huge amount of power compared to the baseline, and for any group of distance managements. Indeed, more than 20% of power reduction is shown.

The second main trend of points, zoomed in Fig 11-c, are dominated by the solutions $12NA/0A/20T$. These solutions do not exceed 1.7% output errors while pushing forward the power consumption reduction. The last trend of solutions, that do not exceed the tolerable output error of 10%, is zoomed in Fig 11-d. These solutions are dominated by the $8NA/4A/20T$ with $BER_{ax} = 10^{-3}$. These solutions provide less than 8.15% of output error while providing the more power consumption reduction.

To conclude this exploration, first we can state that the design space is huge, and designing an ONI that is be able to configure all of any configuration is unrealistic. Secondly, we highlighted that the different solutions can be grouped in four trends, almost no error (Fig 11-b), few errors (Fig 11-c), reasonable errors (Fig 11-d), and unacceptable errors. All trends are dominated by one or few optimal points, hence limiting the number of interesting solutions to embed in the ONIs./JAEI can not understand this sentence As already mentioned, we target to limit the complexity of the ONI and its management by only considering one low power communication scheme. Thus, the most relevant solution to provide the highest power savings is $8NA/4A/20T$ with $BER_{ax} = 10^{-3}$ which provide a reasonable output error. Finally, regarding the distance of communication, the proposed solution of only considering two group of distance (i.e. Short/Long) provides efficient result compared to no distance management technique. Compared to the optimal solution induced by fully distance proportional links, our solution consumes only 12% additional power. However, it is important to notice that our ONI design is realistic as it requires only four laser powers, compared to [13] for which the complexity linearly increases with both number of cores and number of BER levels.

### 6.4 Laser Driver Design Complexity

As discussed in the previous section, fine grain tuning of the optical signal power, as proposed in [13], allows reducing the laser power consumption. However, such approach leads to significant challenges related to the design of both analog and digital electronic circuits to control the lasers. In [13], the authors consider a given approximation level for each application. This implies that each driver is capable of targeting, for each destination, any BER from $10^{-12}$ to $10^{-2}$. Hence, for a N-cluster architecture, each driver is capable of emitting a signal at $11 \times (N-1)$ different power levels (e.g. 165 levels per laser for 16 clusters). Since no hardware evaluation is carried out in [13], we study their driver complexity assuming only two BER values of $10^{-12}$ and $10^{-3}$ for each destination, while still assuming that the optimal BER can be used. Results are provided in Fig. 12 for a number of clusters ranging from $2 \times 2$ to $16 \times 16$. As seen on the figure, the number of power levels for our architecture is constant (4 levels corresponding to $P_{Off}$, $P_H$, $P_M$ and $P_L$) while it linearly grows for *Lorax* despite the considered optimistic assumption. For instance, *Lorax* leads to 255 power levels for 128 clusters (1 level for $P_{Off}$ and two BER levels for each of the 127 destinations), which significantly increases the design complexity of the laser driver, as discussed in the following.

We now study the complexity of the CMOS drivers assuming an OOK modulation for both Lorax and our method. CMOS drivers are key components to deliver the right current to the on-chip lasers according to the selected optical power level. We designed multi-level laser drivers for our approach and *Lorax* assuming 16 clusters, as represented in Fig. 13. The drivers operate as follow: transistor $T_1$ drive a $I_{bias}$ current to maintain the laser near the threshold area and transistors drove by $En_i$ signals define the current in the laser emission area. In case communication is required or is predicted, Signal $/Sleep$ is set to 1, allowing to start driving current and to prepare laser for emission. This allows

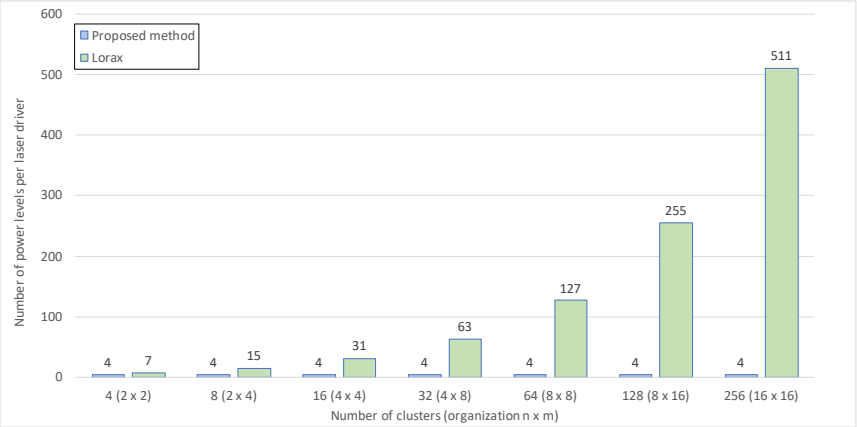


Fig. 12. Number of power levels per laser driver according to the network size for Lorax and the proposed method.

reducing latency penalty occurring to switch on laser from a fully off state [23, 49]. In case no optical communication occurs or is predicted, Signal $/Sleep$ is set to 0, thus leading to energy saving. Temporal penalty and optimization of driver sleep time [50] are out of the scope of this work but would lead to the same energy reduction for both *Lorax* and our method. Input signal D, which corresponds to the data to transmit, controls the laser emission through transistors $T_2$ and $T_3$. The number of transistors increases with the number of laser output powers and the $En_i$ are activated according to selected output power. As detailed in Tables of Fig. 13-c) and d), the more transistors driving current, the higher the laser output power. As a result, the driver for *Lorax* requires 34 transistors to generate the 31 laser power levels (15 destinations and 2 levels per destination, and $P_{Off}$). Our approach requires 7 transistors to generated the 4 power levels (i.e. $P_H$, $P_M$, $P_L$ and $P_{Off}$), hence leading to 4.8× reduction in transistor count.

Table 5 summarizes driver estimations on leakage current and footprint for our proposed method and *Lorax*. When targeting a 65$nm$ technology node from STM and estimation methods from [51–53], the leakage current for our method is 15$nA$ against 125$nA$ for *Lorax* (8× reduction). Regarding the footprint estimations, *Lorax* driver requires 8.89$mm^2$, while our method only requires 0.25$mm^2$ (35× reduction).

| | Number of Power levels | Transistor count | Leakage($nA$) | Footprint($mm^2$) |
|---|---|---|---|---|
| Proposed method | 4 | 7 | 15.65 | 0.253 |
| Lorax [13] | 34 | 32 | 125.12 | 8.89 |
| **Lorax Increase (×)** | **7.75** | **4.85** | **7.99** | **35.14** |

Table 5. Driver area estimations for a 16 cluster architecture

Input bits $En_i$ of the drivers are defined using a digital controller. The controller complexity directly depends on number of configurable laser power levels: while the 4-level driver used for our method requires a 2-input to 3-output controller, Lorax would require a 5-input to 30-output controller. In order to evaluate the area and power overhead, we designed the controllers at the RTL level using VHDL. The controllers were synthesized for a 65$nm$ ST Microelectronics technology using Synopsis Design Compiler. We targeted a 1GHz frequency corresponding to core speed assumed in

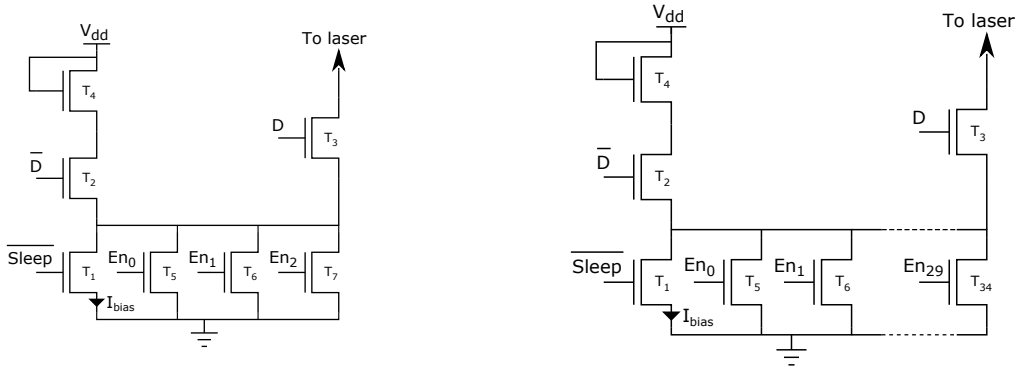


| $En_2$ | $En_1$ | $En_0$ | Current level | Power level |
|---|---|---|---|---|
| 0 | 0 | 0 | $I_{bias}$ | $P_{off}$ |
| 0 | 0 | 1 | $I_L$ | $P_L$ |
| 0 | 1 | 1 | $I_M$ | $P_M$ |
| 1 | 1 | 1 | $I_H$ | $P_H$ |

(c)

| $En_{29}$ | $En_{28}$ | ... | $En_1$ | $En_0$ | Current level | Power level |
|---|---|---|---|---|---|---|
| 0 | 0 | ... | 0 | 0 | $I_{bias}$ | $P_{off}$ |
| 0 | 0 | ... | 0 | 1 | $I_1$ | $P_{ax}$ Hop 1 |
| 0 | 0 | ... | 1 | 1 | $I_2$ | $P_{notax}$ Hop 1 |
| ... | ... | ... | ... | ... | ... | ... |
| 0 | 1 | ... | 1 | 1 | $I_{29}$ | $P_{ax}$ Hop 15 |
| 1 | 1 | ... | 1 | 1 | $I_{30}$ | $P_{notax}$ Hop 15 |

(d)

Fig. 13. Laser drivers assuming 16 clusters for a) proposed approach (4 power levels corresponding to $P_{off}$, $P_H$, $P_M$ and $P_L$) and b) *Lorax* (31 power levels corresponding to $P_{off}$ and 15 destinations with 2 levels per destination). The combinations of inputs signals $En_i$ and the corresponding power level are given in tables c) and d).

the previous experiment. Prime Time was used for power estimation. As seen Table 6, our method leads to controller with an area of $42\mu m^2$, thus leading to a 10 × reduction compared to Lorax ($445\mu m^2$). Similar gain is observed for the power consumption: the method we propose leads a controller consuming a total of $28.58\mu W$ (leakage $3.25\mu W$ and dynamic power $25.33\mu W$) compared to $290\mu W$ for *Lorax*.

| | **Power** | | | **Area** |
|---|---|---|---|---|
| | Dynamic Power ($\mu W$) | Cell Leakage Power ($\mu W$) | Total power($\mu W$) | Area ($\mu m^2$) |
| Proposed method | 25.33 | 3.25 | 28.58 | 42.6 |
| Lorax [13] | 259.8 | 30.9 | 290.7 | 445.1 |
| **Gain** | **10.2** × | **9.5** × | **10.1** × | **10.4** × |

Table 6. Driver controller synthesis results on 65nm STM technology node

For a comprehensive comparison, we evaluate the total power consumption per laser for *Lorax* and our approach, assuming 16 clusters, BER of $10^{-3}$ for approximate communications and 33% laser efficiency [48]. Fig. 14 presents the power breakdown (controller and laser) according to the communication distance. Our approach leads to power consumption of $368\mu W$ and $882\mu W$ for short-range communications (hop count 1 to 5) and long-range (hop count 6 to

16) communications respectively. Regarding the power consumption for *Lorax*, it quadratically increases and leads to average power of i) $573\mu W$ for hop count 1 to 5 and ii) $873\mu W$ for hop count 6 to 16. Hence, our method allows reducing power by 35% for short range communications while leading to only 2.5% power overhead for long range communications. Interestingly, our approach is more energy efficient for the longest communication range (hop count 12 to 15 in the figure). The significant gains obtained for short range communications is possible due to the lower static power of the driver controller. The slight power overhead for long-range communication is acceptable considering the 10× reduction in the area footprint: while the drivers for *Lorax* would occupy a total area of $8.89mm^2$ for the studied 16 clusters architecture, our approach would take an area of $0.25mm^2$.

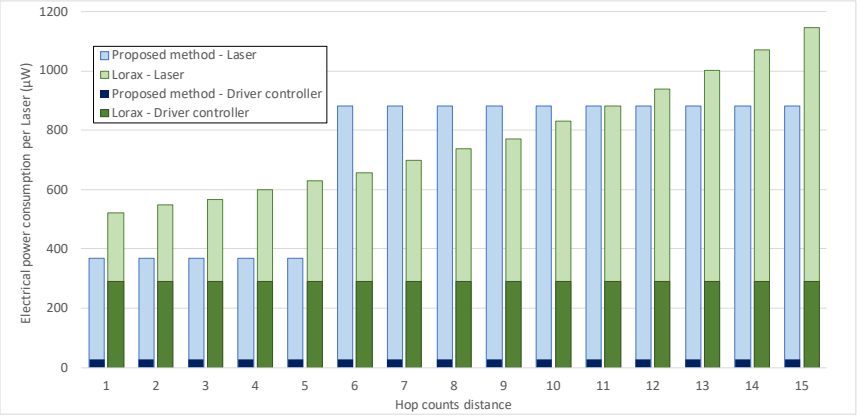


Fig. 14. Total power consumption per hop count for Lorax and the proposed method. Results are given for a 16 clusters architecture assuming $10^{-3}$ BER for approximate communications.

## 7 DISCUSSIONS AND PERSPECTIVES

This study demonstrates the poor scalability of methods involving fine grain tuning of laser powers. It also validates the effectiveness of our approach to maintain the driver complexity low while reaching high level of energy efficiency, especially for short-range and very long-range communications. This opens new research perspectives related to i) short-term deployment of approximate nanophotonic interconnects with mature optical devices and ii) smart mapping of tasks to maximize the utilization of most energy efficient links.

i) Off-chip laser

This paper focuses on on-chip lasers, which are disruptive but promising technologies for large scale integration of light sources on a chip. However, as off-chip laser is more realistic approach at current development of optical chip laser, we will investigate the deployment of the approach to off-chip lasers-based interconnects for our proposal. A key challenge with off-chip Lasers would be to adapt the power level according to the communication needs, as off-chip lasers cannot be dedicated to each IP and are generally shared among them. Usually, the power is configured to satisfy any transmission requirements, hence leading to worst case scenario and significant waste of energy. Approaches involving on-chip optical amplifiers such as [33], or method adapting the power regarding the number of active paths in a tree-based optical NoC [6], could be used to adapt the output power according to the requirement but would significantly increase the design complexity. Furthermore, approaches involving off-chip

lasers necessary require dedicated modulators, while direct modulation could be achieved using on-chip lasers. In this paper, our goal is to evaluate the potential of a prospective architecture to maximize the energy efficiency while keeping the design realistic. To apply the proposed method in the context of off-chip lasers, a straightforward solution would be to consider spatial division multiplexing by using multiple SWMR links per core. Each SWMR would then have the ability to transmit signals at a given power level ($P_H$, $P_M$ and $P_L$) to a specific set of receivers (e.g. SWMR with $P_L$ would reach only receivers part of short-range communications). This would lead to three sets of off-chip laser(s) dedicated to i) approximate / short range communications only (SWMR with $P_L$), ii) accurate / long range communications (SWMR with $P_H$) and iii) both approximate / long range and accurate / short range communications (SWMR with $P_M$). Switching ON or OFF the lasers would be achieved at software level, depending on the executed application. Such approach would ideally complement existing solutions that enable to use only a set of wavelengths to communicate [54] and that could be easily extended to transmitted truncated data. Obviously, such solutions would lead to significant overhead in the design of the photonics chips (number of lasers and layout) as well as on the electronics, which will have to be evaluated.

ii) Task mapping optimization
Optimizing the task mapping in manycore architectures has been thoroughly studied in the context of electrical NoCs [55–57] , hybrid electrical-wireless NoCs [58] and optical NoCs [59]. Since the purpose of existing approaches is mainly to minimize the communication costs, they could be adapted to efficiently use the laser power levels available in our architectures. For instance, it would be possible to prioritize very long-range communications in case $P_H$ is used, which can be obtained by optimizing at system level metrics such pJ/bit/mm. Task mapping will also allow to evaluate the efficiency of the approach when extended to larger set of communication ranges and when applied to application specific domain. For this purpose, Convolutional Neural Networks (CNNs) are particularly interesting candidates as communication volumes and data accuracy requirement strongly change between the CNN layers.

## 8 CONCLUSION

In this paper, we proposed a method to design scalable approximate nanophotonic interconnects. Our approach aims at improving the interconnect energy efficiency by adapting the transmission robustness to the application requirements. We proposed to use both truncation and bits approximation by using lower power signals, thus leading to accurate and approximate communication. We classified communications according to the transmission losses, thus leading to short-range and long range communications. To minimize the number of laser power levels, we also proposed a method allowing to use a same power for both accurate short-range and approximate long-range communications. The method requires only minor modification of the interconnects interface and is scalable due to the reduced number of power levels. We have evaluated the impact of various combination of truncation and bits approximation on the output error of benchmarks applications, namely Blackscholes, Canneal, Fluidanminate and Streamcluster. Results have showed that applications such Streamcluster are the best candidates to use approximation in the nanophotonic interconnect since they are error-resilient and they involve numerous floating-point data. Our method leads to 53% power reduction for an acceptable 8% output errors. Compared to state of the art method involving fine grain laser power tuning, we obtain a 35% power reduction for short-range communication. We also demonstrate 35× and 10× reductions in the laser driver and laser controller sizes respectively. Finally, we shown the opportunities to extend this work to off-chip lasers and to

task mapping optimization, with the aim to take benefits from the significant energy saving available for short-range communications.

## REFERENCES


[1] A. Karkar, T. Mak, K. Tong, and A. Yakovlev, "A survey of emerging interconnects for on-chip efficient multicast and broadcast in many-cores," *IEEE Circuits and Systems Magazine*, vol. 16, no. 1, pp. 58–72, 2016.

[2] W. Wolf, A. A. Jerraya, and G. Martin, "Multiprocessor system-on-chip (mpsoc) technology," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 27, no. 10, pp. 1701–1713, 2008.

[3] T. Alexoudi, N. Terzenidis, S. Pitris, M. Moralis-Pegios, P. Maniotis, C. Vagionas, C. Mitsolidou, G. Mourgias-Alexandris, G. T. Kanellos, A. Miliou, K. Vyrsokinos, and N. Pleros, "Optics in computing: From photonic network-on-chip to chip-to-chip interconnects and disintegrated architectures," *Journal of Lightwave Technology*, vol. 37, pp. 363–379, Jan 2019.

[4] P. K. Hamedani, N. E. Jerger, and S. Hessabi, "Qut: A low-power optical network-on-chip," in *2014 Eighth IEEE/ACM International Symposium on Networks-on-Chip (NoCS)*, pp. 80–87, 2014.

[5] C. Killian, D. Chillet, S. L. Beux, V. Pham, O. Sentieys, and I. O'Connor, "Energy and performance trade-off in nanophotonic interconnects using coding techniques," in *Proceedings of the 54th Annual Design Automation Conference, DAC 2017, Austin, TX, USA, June 18-22, 2017*, pp. 85:1–85:6, ACM, 2017.

[6] L. Zhou and A. K. Kodi, "Probe: Prediction-based optical bandwidth scaling for energy-efficient nocs," in *2013 Seventh IEEE/ACM International Symposium on Networks-on-Chip (NoCS)*, pp. 1–8, April 2013.

[7] J. Postman, T. Krishna, C. Edmonds, L. Peh, and P. Chiang, "Swift: A low-power network-on-chip implementing the token flow control router architecture with swing-reduced interconnects," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 21, no. 8, pp. 1432–1446, 2013.

[8] J. O. Sosa, C. Killian, H. B. Ammar, and D. Chillet, "Min/max time limits and energy penalty of communication scheduling in ring-based ONoC," in *NoCArc 2020*, (On-line, France), Oct. 2020.

[9] Q. Xu, T. Mytkowicz, and N. S. Kim, "Approximate computing: A survey," *IEEE Design Test*, vol. 33, pp. 8–22, Feb 2016.

[10] S. Mittal, "A survey of techniques for approximate computing," *ACM Comput. Surv.*, vol. 48, Mar. 2016.

[11] A. Sampson, W. Dietl, E. Fortuna, D. Gnanapragasam, L. Ceze, and D. Grossman, "Accept: a programmer-guided compiler framework for practical approximate computing," *University of Washington Technical Report UW-CSE-15-01*, vol. 1, June 2015.

[12] J. Lee, C. Killian, S. L. Beux, and D. Chillet, "Approximate nanophotonic interconnects," in *Proceedings of the 13th IEEE/ACM International Symposium on Networks-on-Chip*, NOCS '19, (New York, NY, USA), Association for Computing Machinery, 2019.

[13] F. Sunny, A. Mirza, I. G. Thakkar, S. Pasricha, and M. Nikdast, "LORAX: loss-aware approximations for energy-efficient silicon photonic networks-on-chip," *CoRR*, vol. abs/2002.11289, 2020.

[14] D. Mahajan, K. Ramkrishnan, R. Jariwala, A. Yazdanbakhsh, J. Park, B. Thwaites, A. Nagendrakumar, A. Rahimi, H. Esmaeilzadeh, and K. Bazargan, "Axilog: Abstractions for approximate hardware design and reuse," *IEEE Micro*, vol. 35, no. 5, pp. 16–30, 2015.

[15] A. Sampson, W. Dietl, E. Fortuna, D. Gnanapragasam, L. Ceze, and D. Grossman, "Enerj: Approximate data types for safe and general low-power computation," *SIGPLAN Not.*, vol. 46, pp. 164–174, June 2011.

[16] T. E. Carlson, W. Heirman, S. Eyerman, I. Hur, and L. Eeckhout, "An evaluation of high-level mechanistic core models," *ACM Transactions on Architecture and Code Optimization (TACO)*, 2014.

[17] Y.-H. Kao and H. J. Chao, "Design of a bufferless photonic clos network-on-chip architecture," *IEEE Transactions on Computers*, vol. 63, no. 3, pp. 764–776, 2014.

[18] C. Li, M. Browning, P. V. Gratz, and S. Palermo, "Luminoc: A power-efficient, high-performance, photonic network-on-chip," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 33, no. 6, pp. 826–838, 2014.

[19] C. Chen, J. L. Abellán, and A. Joshi, "Managing laser power in silicon-photonic noc through cache and noc reconfiguration," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 34, no. 6, pp. 972–985, 2015.

[20] B. Neel, M. Kennedy, and A. Kodi, "Dynamic power reduction techniques in on-chip photonic interconnects," in *Proceedings of the 25th Edition on Great Lakes Symposium on VLSI*, GLSVLSI '15, (New York, NY, USA), p. 249–252, Association for Computing Machinery, 2015.

[21] J. Bashir and S. R. Sarangi, "Predict, share, and recycle your way to low-power nanophotonic networks," *J. Emerg. Technol. Comput. Syst.*, vol. 16, Oct. 2019.

[22] S. Van Winkle, A. K. Kodi, R. Bunescu, and A. Louri, "Extending the power-efficiency and performance of photonic interconnects for heterogeneous multicores with machine learning," in *2018 IEEE International Symposium on High Performance Computer Architecture (HPCA)*, pp. 480–491, 2018.

[23] C. Chen and A. Joshi, "Runtime management of laser power in silicon-photonic multibus noc architecture," *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 19, no. 2, pp. 3700713–3700713, 2013.

[24] S. S. Vatsavai, V. S. P. Karempudi, and I. G. Thakkar, "PROTEUS: rule-based self-adaptation in photonic nocs for loss-aware co-management of laser power and performance," in *14th IEEE/ACM International Symposium on Networks-on-Chip, NOCS 2020, Hamburg, Germany, September 24-25, 2020*, pp. 1–8, IEEE, 2020.

[25] R. Boyapati, J. Huang, P. Majumder, K. H. Yum, and E. J. Kim, "Approx-noc: A data approximation framework for network-on-chip architectures," in *2017 ACM/IEEE 44th Annual International Symposium on Computer Architecture (ISCA)*, pp. 666–677, 2017.

[26] L. Wang, X. Wang, and Y. Wang, "Abdtr: Approximation-based dynamic traffic regulation for networks-on-chip systems," in *2017 IEEE International Conference on Computer Design (ICCD)*, pp. 153–160, 2017.

[27] Y. Chen, M. F. Reza, and A. Louri, "Dec-noc: An approximate framework based on dynamic error control with applications to energy-efficient nocs," in *2018 IEEE 36th International Conference on Computer Design (ICCD)*, pp. 480–487, 2018.

[28] M. F. Reza and P. Ampadu, "Approximate communication strategies for energy-efficient and high performance noc: Opportunities and challenges," in *Proceedings of the 2019 on Great Lakes Symposium on VLSI*, (New York, NY, USA), Association for Computing Machinery, 2019.

[29] F. Betzel, K. Khatamifard, H. Suresh, D. J. Lilja, J. Sartori, and U. Karpuzcu, "Approximate communication: Techniques for reducing communication bottlenecks in large-scale parallel systems," *ACM Comput. Surv.*, vol. 51, Jan. 2018.

[30] A. B. Ahmed, D. Fujiki, H. Matsutani, M. Koibuchi, and H. Amano, "Axnoc: Low-power approximate network-on-chips using critical-path isolation," in *2018 Twelfth IEEE/ACM International Symposium on Networks-on-Chip (NOCS)*, (New York, NY, USA), pp. 1–8, IEEE, Oct 2018.

[31] D. S. Khudia, B. Zamirai, M. Samadi, and S. Mahlke, "Rumba: An online quality management system for approximate computing," in *2015 ACM/IEEE 42nd Annual International Symposium on Computer Architecture (ISCA)*, pp. 554–566, 2015.

[32] C. Sun, C. O. Chen, G. Kurian, L. Wei, J. Miller, A. Agarwal, L. Peh, and V. Stojanovic, "Dsent - a tool connecting emerging photonics with electronics for opto-electronic networks-on-chip modeling," in *2012 IEEE/ACM Sixth International Symposium on Networks-on-Chip*, pp. 201–210, May 2012.

[33] I. G. Thakkar, S. V. R. Chittamuru, and S. Pasricha, "Run-time laser power management in photonic nocs with on-chip semiconductor optical amplifiers," *2016 Tenth IEEE/ACM International Symposium on Networks-on-Chip (NOCS)*, pp. 1–4, 2016.

[34] Z. Wang, H. Gu, Y. Yang, and B. Zhang, "Power allocation method for tdm-based optical network on chip," *IEEE Photonics Technology Letters*, vol. 25, no. 10, pp. 973–976, 2013.

[35] F. Mandorlo *et al.*, "Controlled multi-wavelength emission in full cmos compatible micro-lasers for on chip interconnections," *Journal of Lightwave Technology*, vol. 30, pp. 3073–3080, Oct 2012.

[36] Z. Wang a *et al.*, "A holistic modeling and analysis of optical-electrical interfaces for inter/intra-chip interconnects," *IEEE TVLSI*, vol. 24, pp. 2462–2474, July 2016.

[37] B. Radi, M. S. Nezami, M. Taherzadeh-Sani, F. Nabki, M. Ménard, and O. Liboiron-Ladouceur, "A 22-gb/s time-interleaved low-power optical receiver with a two-bit integrating front end," *IEEE Journal of Solid-State Circuits*, vol. 56, no. 1, pp. 310–323, 2021.

[38] M. Bahadori, S. Rumley, H. Jayatilleka, K. Murray, N. A. F. Jaeger, L. Chrostowski, S. Shekhar, and K. Bergman, "Crosstalk penalty in microring-based silicon photonic interconnect systems," *J. Lightwave Technol.*, vol. 34, pp. 4043–4052, Sep 2016.

[39] Y. Solihin, *Fundamentals of Parallel Multicore Architecture.* Chapman & Hall/CRC, 1st ed., 11 2015.

[40] J. Luo, V.-D. Pham, C. Killian, D. Chillet, I. O'Connor, O. Sentieys, and S. Le Beux, "Run-time management of energy-performance trade-off in optical network-on-chip," in *2018 Conference on Design of Circuits and Integrated Systems (DCIS)*, pp. 1–6, 2018.

[41] T.-J. Kao and A. Louri, "Optical multilevel signaling for high bandwidth and power-efficient on-chip interconnects," *IEEE Photonics Technology Letters*, vol. 27, no. 19, pp. 2051–2054, 2015.

[42] P. Dong, W. Qian, S. Liao, H. Liang, C. Kung, N. Feng, R. Shafiiha, J. Fong, D. Feng, A. V. Krishnamoorthy, and M. Asghari, "Low loss silicon waveguides for application of optical interconnects," in *IEEE Photonics Society Summer Topicals 2010*, pp. 191–192, 2010.

[43] M. Bahadori, S. Rumley, D. Nikolova, and K. Bergman, "Comprehensive design space exploration of silicon photonic interconnects," *Journal of Lightwave Technology*, vol. 34, no. 12, pp. 2975–2987, 2016.

[44] N. M. Fahrenkopf, C. McDonough, G. L. Leake, Z. Su, E. Timurdogan, and D. D. Coolbaugh, "The aim photonics mpw: A highly accessible cutting edge technology for rapid prototyping of photonic integrated circuits," *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 25, no. 5, pp. 1–6, 2019.

[45] P. Dong, W. Qian, S. Liao, H. Liang, C.-C. Kung, N.-N. Feng, R. Shafiiha, J. Fong, D. Feng, A. V. Krishnamoorthy, and M. Asghari, "Low loss shallow-ridge silicon waveguides," *Opt. Express*, vol. 18, pp. 14474–14479, Jul 2010.

[46] Q. Wilmart, H. El Dirani, N. Tyler, D. Fowler, S. Malhouitre, S. Garcia, M. Casale, S. Kerdiles, K. Hassan, C. Monat, X. Letartre, A. Kamel, M. Pu, K. Yvind, L. K. Oxenløwe, W. Rabaud, C. Sciancalepore, B. Szelag, and S. Olivier, "A versatile silicon-silicon nitride photonics platform for enhanced functionalities and applications," *Applied Sciences*, vol. 9, no. 2, 2019.

[47] F. Xia, M. Rooks, L. Sekaric, and Y. Vlasov, "Ultra-compact high order ring resonator filters using submicron silicon photonic wires for on-chip optical interconnects," *Opt. Express*, vol. 15, pp. 11934–11941, Sep 2007.

[48] X. Chen, J. Feng, J. Xu, J. Zhang, and S. Chen, "Simultaneously tolerate thermal and process variations through indirect feedback tuning for silicon photonic networks," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 40, no. 7, pp. 1409–1422, 2021.

[49] M. T. Hosseinabadi and N. Ansari, "Multi-power-level energy saving management for passive optical networks," *IEEE/OSA Journal of Optical Communications and Networking*, vol. 6, no. 11, pp. 965–973, 2014.

[50] F. Lan, R. Wu, C. Zhang, Y. Pan, and K. T. Cheng, "Dlps: Dynamic laser power scaling for optical network-on-chip," in *2017 22nd Asia and South Pacific Design Automation Conference (ASP-DAC)*, pp. 726–731, 2017.

[51] Y. Agrawal, R. Chandel, and R. Dhiman, "Design and analysis of efficient multilevel receiver for current mode interconnect system," in *2014 IEEE Students' Conference on Electrical, Electronics and Computer Science*, pp. 1–6, 2014.

[52] M. Dave, M. Jain, M. Shojaei Baghini, and D. Sharma, "A variation tolerant current-mode signaling scheme for on-chip interconnects," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 21, no. 2, pp. 342–353, 2013.

[53] Y. Agrawal, R. Chandel, and R. Dhiman, "High-performance current mode receiver design for on-chip vlsi interconnects," *Advances in Intelligent Systems and Computing*, vol. 343, pp. 527–536, 01 2015.

[54] A. Narayan, Y. Thonnart, P. Vivet, and A. K. Coskun, "Prowaves: Proactive runtime wavelength selection for energy-efficient photonic nocs," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 40, no. 10, pp. 2156–2169, 2021.

[55] C. Wu, C. Deng, L. Liu, J. Han, J. Chen, S. Yin, and S. Wei, "A multi-objective model oriented mapping approach for noc-based computing systems," *IEEE Transactions on Parallel and Distributed Systems*, vol. 28, no. 3, pp. 662–676, 2017.

[56] X. Wang, T.-M. Choi, X. Yue, M. Zhang, and W. Du, "An effective optimization algorithm for application mapping in network-on-chip designs," *IEEE Transactions on Industrial Electronics*, vol. 67, no. 7, pp. 5798–5809, 2020.

[57] L. Yang, W. Liu, N. Guan, and N. Dutt, "Optimal application mapping and scheduling for network-on-chips with computation in stt-ram based router," *IEEE Transactions on Computers*, vol. 68, no. 8, pp. 1174–1189, 2019.

[58] W. Choi, K. Duraisamy, R. G. Kim, J. R. Doppa, P. P. Pande, D. Marculescu, and R. Marculescu, "On-chip communication network for efficient training of deep convolutional networks on heterogeneous manycore systems," *IEEE Transactions on Computers*, vol. 67, no. 5, pp. 672–686, 2018.

[59] Y. Wang and K.-T. Cheng, "Task mapping-assisted laser power scaling for optical network-on-chips," in *2019 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*, pp. 1–6, 2019.