Online Appendix : Predict, Share, and Recycle your Way to Low Power Nanophotonic Networks

Janibul Bashir and Smruti R. Sarangi
Department of Computer Science and Engineering, Indian Institute of Technology, Delhi - 110016
Email: {janibbashir,srsrarangi}@cse.iitd.ac.in

1 Introduction

Silicon photonics is a very promising technology, because of the inherent advantages of light such as high bandwidth due to wavelength division multiplexing (WDM), high speed (7ps/mm) and the favorable power-performance tradeoffs while building optical networks [4, 7, 16, 38, 43, 46]. Because of these advantages various photonic on-chip networks (PNOC) have been proposed in the literature [5, 6, 18, 45, 70, 73, 95]. Most of these proposals try to solve a major shortcoming of PNOCs – static power consumption.

PShaRe is a PNOC, which incorporates several novel ideas to decrease static power consumption by nearly 4.7X as compared to other state of the art designs. In addition, it has several performance improving features as well. In this technical report we provide some details that could not be added to the main paper because of a lack of space.

The organization of this report is as follows:

Section	Contents
Section 2	Detailed background, feasibility of silicon photonics, and the design of common optical devices, and the overall feasibility of *PShaRe*.
Section 3	Related work.
Section 4	The design of the artificial neural network (ANN) used in laser modulation.
Section 5	The area and power analysis of *PShaRe*.
Section 6	Representative server details.
Section 7	Thermal tunning, Tuning controller and detailed thermal analysis.

2 Background and Feasibility of Silicon Photonics

2.1 Basic Optical Communication

Figure 1 shows the basic optical communication architecture inside the chip. The light from the off-chip laser source is coupled into the chip using special tapered couplers in order to reduce the coupling loss. The laser power (light at 1550nm) then passes through power or backbone waveguides to reach all the optical stations.

Let us suppose station A (optical receiver + transmitter) wants to send data to station B. It diverts a fraction of the transmitted optical power from the backbone waveguide, uses a comb splitter [53] to generate signals at 32-64 equispaced wavelengths, and modulates each wavelength separately using an electrically controlled modulator. It subsequently inserts the modulated signals into the data waveguide. At the receiving end we have an array of filters that individually filter each of the multiplexed wavelengths and pass them to photodetectors. These photodetectors detect the presence of an optical signal at the specified wavelength. We can thus use this mechanism to encode, transmit and subsequently decode digital messages.

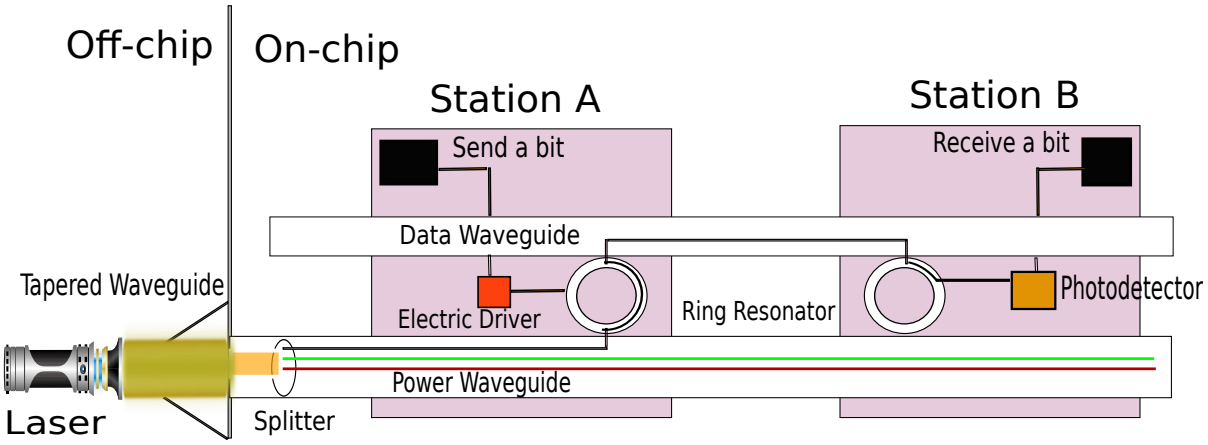

Figure 1: On-chip optical communication

2.2 Feasibility

2.2.1 Feasibility of Silicon Photonics

The device physics and computer architecture communities have been working in the broad silicon photonics area for the last 15 years. There are hundreds of highly cited papers published in this area and these are well accepted in these communities [8, 9, 22, 34, 47–49, 51, 64, 69, 88, 92].

Secondly, as of today many commercial companies have started medium scale manufacturing of systems based on silicon photonics. In 2012, IBM fabricated a 1Tbps parallel interconnect using VCSELs and photodetectors at 90nm technology [3]. In 2015, they designed and tested a fully functional integrated wavelength-multiplexed silicon photonics chip. Moreover, they have a dedicated group working on silicon photonic devices [40]. On similar lines, a silicon photonic startup 'Sicoya' in 2016 developed a fully integrated silicon photonics transceiver chip that is designed for server interconnects [83]. Likewise, HP labs has manufactured robust high speed ring resonators [89], and these are now considered to be mature prototypes.

Some other companies such as Luxtera [60], Mellanox [62] and Acacia communications [1] have started manufacturing silicon photonics devices for on-chip networks. Even in academia, researchers at UC Berkeley have fabricated a 2 core chip with 850 optical components, including ring resonators [67]. All of these are fully functional and robust devices.

2.2.2 Feasibility of Optical Devices

The photonic device parameters that we have used in our architecture are standard parameters. These parameters have been validated in the literature using multiple methods: direct measurements on fabricated chips, device simulations, and theoretical predictions [13, 23, 30, 33, 77, 81, 99]. In addition, multiple groups all over the world have come to the same conclusions regarding the values of the parameters and the resultant performance figures of the fabricated optical devices. In addition, we have done our own experiments while designing and characterizing our device library. All the optical devices such as ring resonators, bends, y-junctions, waveguides and splitters have been rigorously validated by performing simulations on Synopsys RSoft [87] and Lumerical [59, 84].

Lasers *PShaRe* uses a combination of on-chip and off-chip lasers in order to have a power efficient optical on-chip network. Other prior works have also used such combination of lasers [21, 95] to achieve their objectives.

For off-chip lasers we use commercially available directly modulated lasers (DML) (examples: Fitel FOL15DDBA, Finisar DM 80). These lasers have even been fabricated by many researchers and provide very high modulation rates (@ GHz speed) [11, 27–29, 41]. The main issue in using such lasers is their thermal stability. However, it has been addressed to a large extent by Fukamachi et al. [31] . They demonstrated a thermally stable and a much efficient DML laser, showing a stable operation upto 100° with a modulation speed of 25Gbps. Moreover, an array of DML lasers can be used to create a tunable laser with different power layers [74].

For on-chip lasers, we can use any wavelength division multiplexing compatible laser with fast switching time.

Such lasers are mainly built using III-V compound semiconductors, including the Ge and InP lasers [12, 15, 24, 25, 44, 50, 58, 71, 78, 80, 86, 86, 100]. Both InP and Ge based lasers have been fabricated [12, 25, 50, 71] and can be easily integrated on the chip. Moreover, it has been shown that these lasers can easily operate at room temperature and are DWDM compatible [12,58]. Thus, such lasers are the best candidates for on-chip laser sources and we have assumed same lasers in our design. Besides InP and Ge lasers, VCSEL lasers can also be used as on-chip lasers [63, 75, 96] because of their very low turn-on time (less than 100ps [75]). However, due to some disadvantages such as emitting significant heat inside the chip [63, 75], we have not considered such lasers in our work.

Waveguides

Waveguides are the optical links which acts as a channel for the optical information to propagate from one location to another. The three major types of waveguides used in on-chip optical communication are shown in Figure 2. However, among these rib waveguide is the most commonly used because it results in less power loss as compared to other waveguides [93], and a s a result *PShaRe* has used same waveguides in its design.

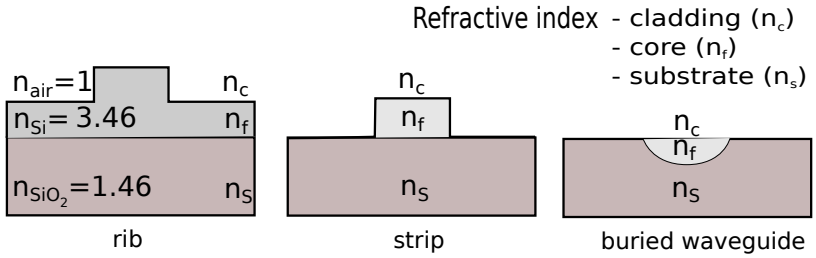

Figure 2: Optical waveguides

Waveguide Bend We require a waveguide bend in op-chip networks in order to implement a bend in the communication channels. We simulated a 90° bend using Synopsys RSoft tool. The full wave simulation is given in Figure 3. The bends are associated with high optical losses and the amount of loss is dependent upton the bend parameters such as rib width, inner and outer radius, and etch depth.

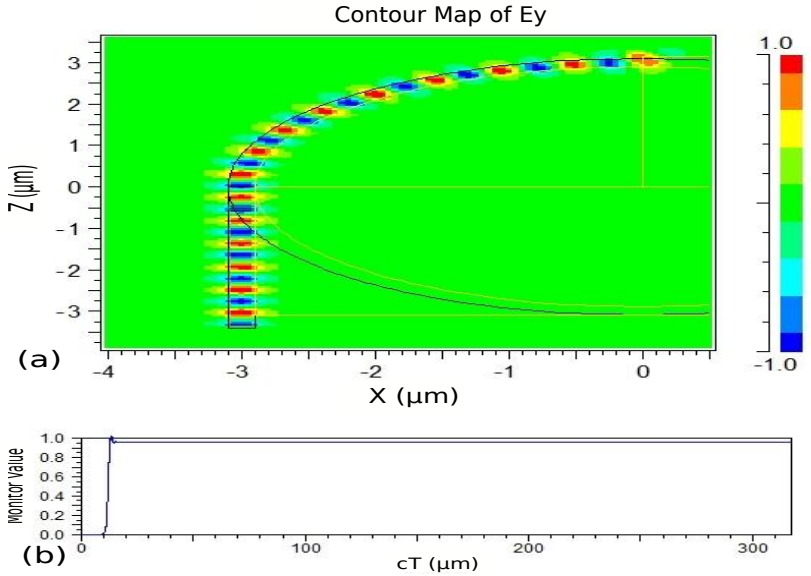

Figure 3: RSoft simulation of a waveguide bend

Y-Junction Y-juntion is an optical component which is used in on-chip networks in order to split or combine the light. Y-junction can be used in symmetrical or assymmetrical mode. In symmetrical mode it divides the light into equal parts wheres in assymmetrical behaviour it splits the light into unequal parts. The RSoft simulation of a Y-junction is given in Figure 4.

Ring Resonators A ring resonator is one main component in the on-chip optical networks. It is used both at the transmitter side and at the receiver sie. At the transmitter side it acts as a modulator and inserts the light into the data waveguide, whereas at the receiver end it acts as a filter. It filters out a specific wavelength form the data waveguide

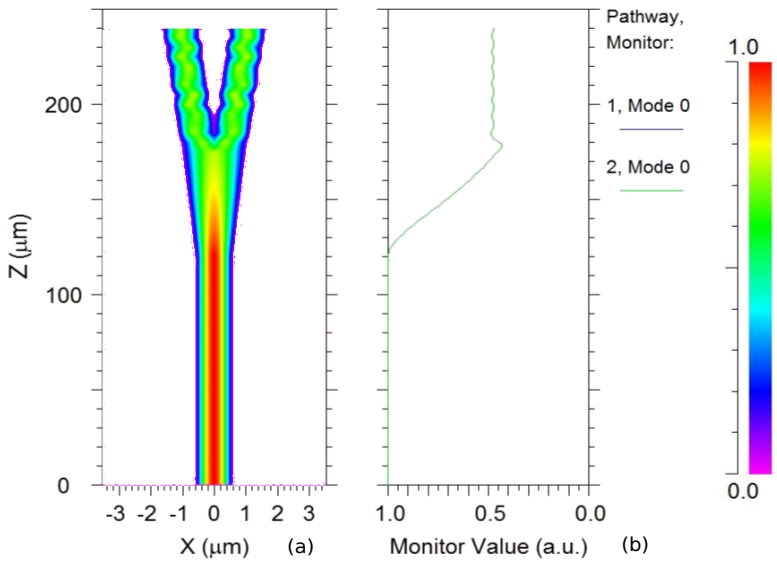

Figure 4: RSoft simulation of a Y junction

and guides it to the photodetector. Figure 5 shows the design of a ring resonator and the Figure 6 gives the RSoft simulation results of a ring resonator.

Many research prototypes have been proposed for ring resonators [26, 52, 57, 99]. Even the researchers at HP have demonstrated and fabricated the ring resonator based transmitter achieving 9Gb/s data rate [54]. In our design, We have used a popular microring resonator proposed by Xu et al. [99], which is area efficient and can modulate data at a rate of up to 10 GHz.

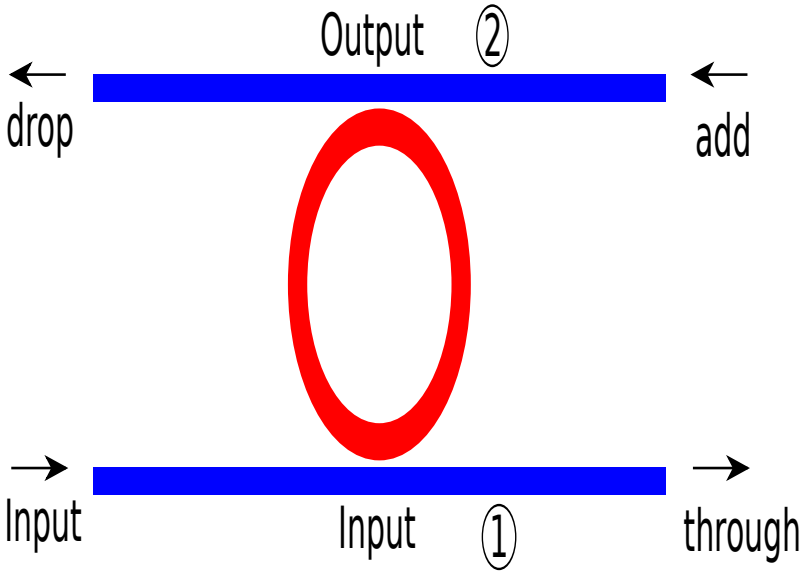

Figure 5: Ring resonator

2.2.3 Feasibility of the System

Let us compare *PShaRe* to similar proposals: Coldbus [74] and Probe [104]. *PShaRe* requires additional area to accommodate the photonic heaters, on-chip lasers, representative servers and ANN hardware. The additional area is roughly $3mm^2$ ($< 1.5\%$ area overhead in a typical 260 mm^2 die). However, it has resulted in reduction in number of ring resonators, and hence we have saved some space by decreasing the counts of optical components (see Table 1). For example, for a 20 node chip, Corona [92] requires 280K ring resonators; however, *PShaRe* requires only 47K ring resonators.

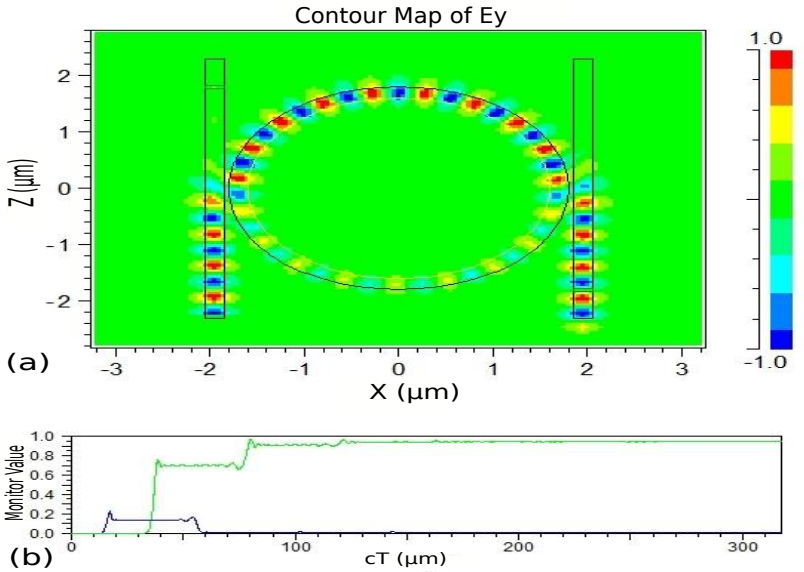

Figure 6: Microring RSoft simulation

Topology	# of waveguides	# of ring resonators
Token-Ring	130	130K
Flexishare [69]	60	70K
ColdBus [74]	56	48K
Corona [92]	132	280K
PShaRe	181	47K

Table 1: Comparison of the number of optical components

To summarize, all the devices and technologies that we use have already been created and have been shown to work for all process corners. We are thus confident about the feasibility of our design. In addition, photonic architectures have already earned a lot of credibility with dozens of commercial optical devices, and hundreds of research prototypes.

3 Related Work

The quintessential method of reducing static power consumption is by turning off the laser (or by reducing its intensity) during periods of inactivity. An easy method is to use on-chip lasers as has been done by Xiaowen et al. [98]. Here, we can place small laser sources near each station and activate them on-demand.

However, on-chip lasers have their shortcomings. They have low wall plug efficiencies and thus all the energy that they draw from the supplies will get dissipated within the chip. This will cause severe thermal and reliability problems. Hence, large scale integration of on-chip lasers is not desirable.

As a result, using off-chip lasers for a bulk of the communication is more common and advisable [69, 70, 74, 92, 104]. The issues here are that we need to predict the laser usage in the future, and then modulate the laser accordingly. Note that it takes time to collate predictions, and to send it to the laser. Hence, we need a method to predict network activity in the future.

3.1 Prediction

We divide time into fixed size durations called *epochs*. At every epoch, we try to predict the laser power requirement for the next few epochs. This information is collected, processed, and transmitted to the laser.

Different proposals use different parameters to predict the amount of laser power that is required. *Probe* [104] predicts the power requirement based on the link and buffer utilization, whereas *ColdBus* [74] predicts based on

the PC addresses of memory instructions. The basic insight that the authors of ColdBus use is that certain memory instructions (identified by their PCs) have a high likelihood of missing in the cache. Their occurrence in the pipeline thus indicates that we might have some network activity resulting out of a cache miss. ESPN [56] divides the network into various sub-networks and uses an independent laser for each sub-network. These lasers are turned on and off independently based on the runtime bandwidth estimate. Similarly, Pulse [65] uses information about cache-line sharers to estimate the amount and nature of network traffic. Most recent work in this area uses a linear regression model to predict the number of packets injected into the network and then accordingly scales the laser power [95].

Almost all such predictors try to capture the linear behaviour of the traffic patterns in the network. However, the traffic patterns in an on-chip network exhibit non-linear relationships, which can be captured easily with the help of non-linear predictors. Thus, we use a very different kind of predictor that uses an artificial neural network to capture the irregularities in the network traffic in order to predict the network activity, and its enhanced accuracy gives us both performance and power benefits.

3.2 Sharing

The other method to decrease the static power consumption is to allow the stations to share the network resources such as waveguides. This decreases the number of waveguides, ring resonators, and consequently the trimming power. For example, Flexishare [69] allows stations to share the available bandwidth and laser power by creating a MWMR network. This results in a decrease in the number of optical channels and at the same time it increases the utilization of laser power. On the same lines, Arslan et al. [105] propose a wavelength stealing method in which a station is allowed to use waveguides belonging to some other station. However, there are chances of collisions. To handle such collisions, it uses erasure coding techniques to recover the messages affected due to collisions. XShare [20] also works on similar lines, where it combines multiple small packets from different stations to form a larger packet, and sends the large packet on the waveguides. The method of sharing requires an arbitration mechanism to provide mutual exclusion. The most commonly used method is the token based arbitration [69] scheme. However, this scheme results in high static power consumption because the tokens are continuously flowing through the token waveguide.

In our design, we allow stations to share the power and data waveguides using a distributed arbitration scheme. Our scheme reduces the static power consumption as compared to the token based scheme because there is no continuous flow of tokens in our design. Finally, to the best of our knowledge, there is no research proposal that has looked at recycling wasted optical power, which is our **novel** contribution.

4 ANN Based Predictor

The main aim of any laser power predictor cum modulator is to scale the off-chip laser power in on-chip optical networks by determining the amount of laser power required in the near future and then modulating the laser accordingly. The standard approach is to divide the execution time into fixed size durations, called epochs. The predictor predicts the laser power requirement for the next epoch and modulates the laser power at the beginning of the next epoch. This will decrease the static power consumption to a large extent. However, in order to accurately predict the laser power requirement, it is necessary to find the relation of the network related parameters with the network traffic. The predictors try to learn this relationship, and determine the laser power requirement based on the set of network related parameters.

4.1 Motivation

All previous proposals have used a set of runtime features, which are assumed to be linearly separable. The predictors used in ColdBus [74] and Probe [104] use the sequence of PC addresses or past network activity to predict the future network activity respectively.

Similarly, machine learning based predictors use a set of architectural parameters such as L2 requests, core requests, core responses, L1 up and down requests and many more to determine the number of packets injected into the network [95]. These ML predictors try to learn a linear relationship.

However, after analyzing the behavior of benchmarks, we concluded that there is a significant degree of non-linearity. Thus, linear predictors will not be able to learn such non-linear relationships, and thus will result in less

accurate predictors.

Moreover, a predictor should use a mix of architectural and network related parameters in order to have a higher prediction accuracy. Such a predictor has hitherto not been designed. As a result we believe that an ANN (artificial neural network) based predictor can easily learn the non-linear relationship for a set of network and architectural parameters and will provide better prediction results.

4.2 Feature Selection

The features selected are used as inputs to the predictor and the predictor predicts whether any laser power is required in the next epoch or not. The main characteristics of the features are: they should be ready available at every optical station, and they should have minimal hardware requirements (if generated). The initial feature set selected for our predictor includes both architectural (AP) and network parameters (NP) as shown in Table 2. The number of features used greatly affects the power and area of the overall system. However, having a large number of features increases the prediction accuracy. Thus, there is a tradeoff between the accuracy and the area/power overheads of the predictor.

1. L1 up requests - L1 to core (AP).
2. L1 down requests - L1 to L2 (AP).
3. Private cache evictions (AP).
4. L2 up requests - L2 to L1(AP).
5. L2 down requests - L2 to lower memory (AP).
6. Requests from a core to the optical station (AP).
7. Packets sent to a core (AP).
8. Packets sent to other optical stations (NP).
9. Requests from other optical stations (NP).
10. Average waiting time (NP).
11. Link utilization (NP).
12. Buffer utilization (NP).
13. Pending requests at the queue (NP).
14. Network history in previous epochs (AP + NP).

Table 2: Inputs to the ANN predictor

By carrying out initial simulations and after analyzing the results, we tried to eliminate the features that barely affect the prediction accuracy. The main aim was to decrease the area and power overheads of the predictor without affecting the prediction accuracy. We tried to eliminate as many features as possible and the most effective set consists of the parameters shown in Table 3. We found these three sets of parameters to be enough.

1. Private Cache Evictions in the current epoch.
2. Pending Events at an optical station.
3. Network traffic injected in the last five epochs.

Table 3: Final feature list

4.3 Evaluation Setup

In order to create an ANN based predictor, we require three datasets: training, validation and testing datasets. The training dataset is used to learn the non-linear relationships across the set of parameters and to fine tune the weights. The validation data set is used tune the various parameters of the predictor. In our case we have used the validation dataset to determine the optimal number of neurons that should be used in the hidden layer. Finally, the testing dataset is used to test the prediction accuracy of our final design.

We used the benchmarks from the Splash 2 [97] and Parsec [10] benchmark suites, and the Tejas [82] simulator to collect the data and create the different datasets. In our simulator, we used various counters to store the value of different parameters in the current epoch. Most of the counters (other than the ones to store the history of the last 5 epochs) are reset at the end of every epoch.

We have a total of 13 benchmarks from the two benchmark suites. These include a combination of memory intensive and compute intensive benchmarks. In order to create a training set, we choose 3 benchmarks from the Parsec and 2 from the Splash suites respectively. These benchmarks are combined to create 6 benchmark pairs. These pairs are run simultaneously in order to train the predictor. For validation, we use 2 benchmarks from each suite to create 4 benchmark pairs. In order to test the final predictor, we ran each benchmark separately. The final results show that the prediction accuracy of our predictor is close to 95%, which is significantly more than other predictors. Thus, it is clear that the ANN based predictor is best choice for predicting the network traffic and consequently scaling the off-chip laser power.

5 Area and Power Estimation

5.1 Area Estimation

Parameter	**Value**	**Parameter**	**Value**
Cores	32	Technology	14nm
Frequency	2.5 GHz		
Processor Core			
pipeline	Four-issue out-of-order	IW size	54
iTLB	128 entry	dTLB	128 entry
Private L1 i-cache, d-cache			
Write-mode	Write-back	Block size	64 bytes
Associativity	4	Size	64 kB
Latency	2 cycles	MSHRs	32
Shared L2			
Write-mode	Write-back	Block size	64 bytes
Associativity	4	# banks	32
Latency (per bank)	8 cycles	Bank size	512 KB
Main Memory			
Latency	150 cycles	Mem. controllers	4
Queue Sizes			
Optical Station Queue	16		
Electrical NoC			
Topology	Flattened Butterfly	Routing Alg.	X-Y
Flit size	16 bytes	Hop-latency	1 cycle
Routing delay	2/3 cycles	# Virt. channels	4
(w/wo bypassing)		Buffers/port	8
Auxiliary structures (size in number of entries)			
RCB	128	VB	20
MQ	16		

Table 4: Simulation parameters (also see [51])

In this section we estimate the area occupied by various electrical and optical components on the chip. Electrical components refer to the components that are present in the chip irrespective of what kind of on-chip network is used, whereas optical components include all those components that are required for the working of a PNOC. The main electrical components are the cores, caches, and memory controllers. The optical components include ring resonators, photonic heaters, on-chip lasers, waveguides, representative servers, and the ANN predictor.

Table 4 gives the architectural parameters of our design and the parameters of an electrical NoC. Table 5 gives the counts of various optical components used in our chip and Table 6 shows the area of each component. The total area occupied by electrical and optical components is shown in Table 7.

Optical system	# of waveguides	# of ring resonators
Power Delivery	64	≈2K
Data Network	64	≈41K
Arbitration	20	≈2K
Prediction	1	≈100
Tuning	32	≈1K
Total	181	≈47K

Table 5: Number of optical components

Structures	Area per structure
Prediction	
Gates for multiplications and additions	$146\mu m^2$ per neuron [2,91]
Lookup table for sigmoid	$780nm^2$ per neuron [61]
Sharing	
Ring resonator	$113\mu m^2$ [57]
Optical waveguide	$0.3mm^2$ [79]
On-chip lasers	$7.68 \times 10^{-3} mm^2$ [95]
Waveguide Reservation Table (WRT)	8 Bytes (calculated using Cacti 6.0 [66])
Power Reuse	
Optical Heaters (OH)	$\approx 768\mu m^2$ [37]
Near Field Transducers (NFT)	$\approx 10000nm^2$ [14]
Tuning Controller	$\approx 4000\mu m^2$

Table 6: Area of each component

Component	Area
Electrical components	
Cores including private caches	$128mm^2$
L2 cache	$60mm^2$
Optical components	
Microring resonators	$18.2mm^2$
Optical waveguides	$54mm^2$
On-chip lasers	$1.5mm^2$
Representative Servers	$0.7mm^2$
Sigmoid lookup table	$0.11\mu m^2$
Neuron Gates	$0.022mm^2$
Photonic heaters	$0.42mm^2$
Near Field Transducers	$0.52\mu m^2$
Tuning Controller	$\approx 4000\mu m^2$

Table 7: Overall area distribution

Let us briefly analyse the area of each component by taking the ARM cortex A15 processor as a reference design. We assume similar cores in our chip with some minor changes. Using the popularly used scaling factors provided by Stan et al. [85], we compute the size of the core to be less than $4mm^2$ at 14nm technology. With such a small core, it is possible to have 32 cores occupying $128mm^2$. We can also integrate 32 cache banks of capacity 16MB with an area less than $60mm^2$ (calculated using Cacti 6.0 [66] and scaled using [85]). For optical components, we calculated the area of each component and all the components occupy less than $60mm^2$ area on the die. If we budget an additional 20% (results from Intels SCC processor) of area for memory controllers, and other additional structures, our total chip area comes to $260mm^2$, which is the size of a standard die for high end processors. **Note that our approach is not specific to the reference design**.

It should be noted that we are using a 2.5D integration technique. In this integration, we have a separate optical network layer over a logical layer. The communication and power distribution across the layers is provided with the help of vertical interconnects. In our design, we are using the popular silicon through-vias as vertical interconnects [35, 72, 90, 101, 102].

5.2 PShaRe Power Model

In this section we will discuss the power model used in *PShaRe*. We calculated the power consumption using this standard model and then compared the power consumption of *PShaRe* with some of the state of the art PNOCs.

The power consumed in any PNOC has two major components – electrical power and photonic power. Electrical power includes the power consumed by electrical components of the network such as drivers, routers, electrical links, and receivers, and the power consumed for electrical to optical and optical to electrical conversions (E/O and O/E). Photonic power includes the power consumed in sending the messages in the network and the power required to tune the microring resonators. Thus, we can say that the total power is the sum of electrical and optical power as given by Equation 1.

Optical parameters	
Wavelength (λ)	$1.55\mu m$
Width of waveguide (W_g)	$3\mu m$
Slab height	$1\mu m$
Rib height	$3\mu m$
Refractive Index of $SiO_2(n_r)$	1.46
Refractive Index of Si (n_c)	3.45
Input Driver Power	76 μW
Insertion Coupling Loss	50%
Photodetector quantum efficiency	0.8 A/W
Photodetector minimum power	36 μW
Combined transmitter and receiver delay	180-270 ps
Optical propagation delay	7 ps/mm
Electrical propagation delay	35 ps/mm
Bending Loss	1 dB
Waveguide Loss	1 dB
Coupler Loss	1 dB
Photodetector	0.1 dB
Wall Plug Efficiency	30 %
Splitter Loss	0.36 dB
Ring Modulation	$1\mu W/^{\circ}C$
Input Photonic heater power	$1mW/^{\circ}C$
Input Electrical heater power	$34.2\mu W/^{\circ}C$

Table 8: Optical parameters [19, 26, 65, 79]

$$T_{power} = E_{power} + P_{power} \tag{1}$$

Here T_{power} denotes the total power, E_{power} is the electrical power, and P_{power} is the optical (photonic) power.

5.2.1 Electric power model

The electric power includes the power consumed by the electrical links (E_{link}), electric routers (E_{router}) , drivers (E_{driver}) , receivers ($E_{receiver}$), ANN predictor ($E_{predictor}$), and power consumed in E/O/E (E_{eoe}) conversion as shown in Equation 2.

$$E_{power} = E_{link} + E_{router} + E_{driver} + E_{receiver} + E_{predictor} + E_{eoe} \tag{2}$$

In *PShaRe* E_{driver}, $E_{predictor}$, and $E_{receiver}$ are the main components of electric power consumption. However, in electric NoCs E_{link} and E_{router} are the major components. E_{eoe} contributes in a big way to E_{power} in the case of hierarchical PNOCs. This is not the case here. In *PShaRe*, E/O/E conversions do not occur frequently and hence do not contribute measurably to the overall power consumption.

For an ENoC considered in our main paper, we have assumed state of the art electrical links between two routers, consuming 13pJ [7,42,70] of energy per link at 14nm and the energy required by a 128-bit flit to traverse the crossbar of a 5×5 electric switch is 16pJ [42, 70]. Thus, a 128-bit flit requires $29pJ/hop/flit$.

In *PShaRe*, one major component in the electric power consumption is the power consumed by the ANN. In each neuron of the hidden layer we perform 7 additions and 7 multiplications and the output layer neuron performs 6 additions and 6 multiplications, making a total of 48 multiplications and 48 additions in 2 ns [95]. This consumes 21pJ [39] of energy. Thus for an epoch of hundred cycles we require at most $150\mu Watts$ of power (including both the training and the prediction phases). The E/O conversions in *PShaRe* occur at the sender and O/E conversion at the receiver. At the sender, it is done with the help of electric drivers and ring resonators and at the receiver photodetectors perform the O/E conversion. The energy consumed by each such conversion (E/O/E) is $50fJ/bit$ [55].

5.2.2 Photonic Power Model

The two main contributors to the photonic power consumption in PNOCs are the laser power (P_{laser}) and the tuning power (P_{tuning}). Thus, we can write:

$$P_{power} = P_{laser} + P_{tuning} \tag{3}$$

Laser Power: The laser power includes the optical power consumed in sending the messages across different nodes in an on-chip network. It includes the insertion losses (I_{loss}) that occur during the propagation of messages through the optical waveguides. The insertion loss is the sum of various optical losses such as propagation loss through a waveguide (P_{loss}), splitter loss (S_{loss}), bending loss (B_{loss}), photodetecor loss (Pd_{loss}), ring resonator loss (R_{loss}). Please note that the ring insertion and through loss are almost equal and as a result we have included both of these in a single term.

Let us now describe the standard method that we have used in our paper to compute the laser power consumption. We first consider the optical power required by the farthest receiver to detect the message (let this be P_{th}). After this we calculate the minimum power required at the modulator (P_{mod}) by calculating backwards from the receiver to the sender. We factor in all the losses by taking the values from Table 8 . The final laser power is calculated using Equation 4.

$$P_{mod} = P_{th} \times T_c \times WG_{loss} \times RI_{loss} \times Other_{loss}. \tag{4}$$

Here,

P_{th} is the minimum power required by the photodetector to detect a message and is equal to $36\mu W$.

T_c is the clock cycle time.

$WG_{loss} = 10^{L_{wg} \times P_{loss}/10}$ is the waveguide propagation loss. L_{wg} is the length of the waveguide and W_{loss} is the propagation loss in dB/cm.

$RI_{loss} = 10^{N_{ring} \times R_{loss}}$ is the optical power loss during the insertion and passing of light through the ring resonators. N_{ring} is the number of ring resonators attached to a waveguide and R_{loss} is the ring through loss (dB/ring).

$Other_{loss} = 10^{(S_{loss}+B_{loss}+Pd_{loss})/10}$, which incorporates all the other losses that occur during the message's propagation through the on-chip optical network.

The Equation 4 can be re-written as:

$$P_{mod} = P_{th} \times T_c \times 10^{L_{wg} \times P_{loss}/10} \times 10^{N_{ring} \times R_{loss}} \times 10^{(S_{loss}+B_{loss}+Pd_{loss})/10}. \tag{5}$$

Using the values from Table 8 in Equation 5, we calculated the average energy required to send a single bit in *PShaRe*. It is equal to 1pJ (similar values given in [17, 76]).

Tuning Power: In any PNOC, ring resonators are the major building blocks [8, 32, 52]. They are used for different purposes such as modulators, filters and even power splitters [33]. However, ring resonators have one major drawback – temperature dependent resonant wavelength. As a result whenever there is a change in the chip temperature, it results in a change in the resonant wavelength of the ring resonator, which may result in errors. However, thermal tuning is one of the solutions to tune the ring resonators. Here, we set the temperature of ring resonators to a pre-specified temperature. This is done using micro-heaters.

The power required to tune the ring resonators is called the tuning power and it is one of the major components in the total power consumption. In our simulations, we have assumed $1\mu W$ per ring per °C [36, 68–70, 94]. Thus, considering a 20K temperature window, the maximum tuning power required is $20\mu W$ per ring. However, in *PShaRe* we reuse unused optical power, and as a result we decreased the temperature window by nearly 7.6K. This led to overall power savings.

6 Representative Server Details

To control the sharing of optical waveguides between the optical stations , we propose to use a single representative server (RS) for a group of 5 optical stations. The RSs are physically placed at the center of the chip. This cluster containing RSs is called Representative cluster. The representative cluster is associated with a separate optical station, called *Server Station*. This optical station is required to facilitate the communication of representative servers with the remaining optical stations in the optical network. The *Server Station* is connected to the other optical stations using separate waveguides called reservation waveguides. In order to avoid waveguide crossings,

the reservation waveguides are assumed running parallel to the data and power waveguides in the serpentine shaped layout as shown in Figure 7.

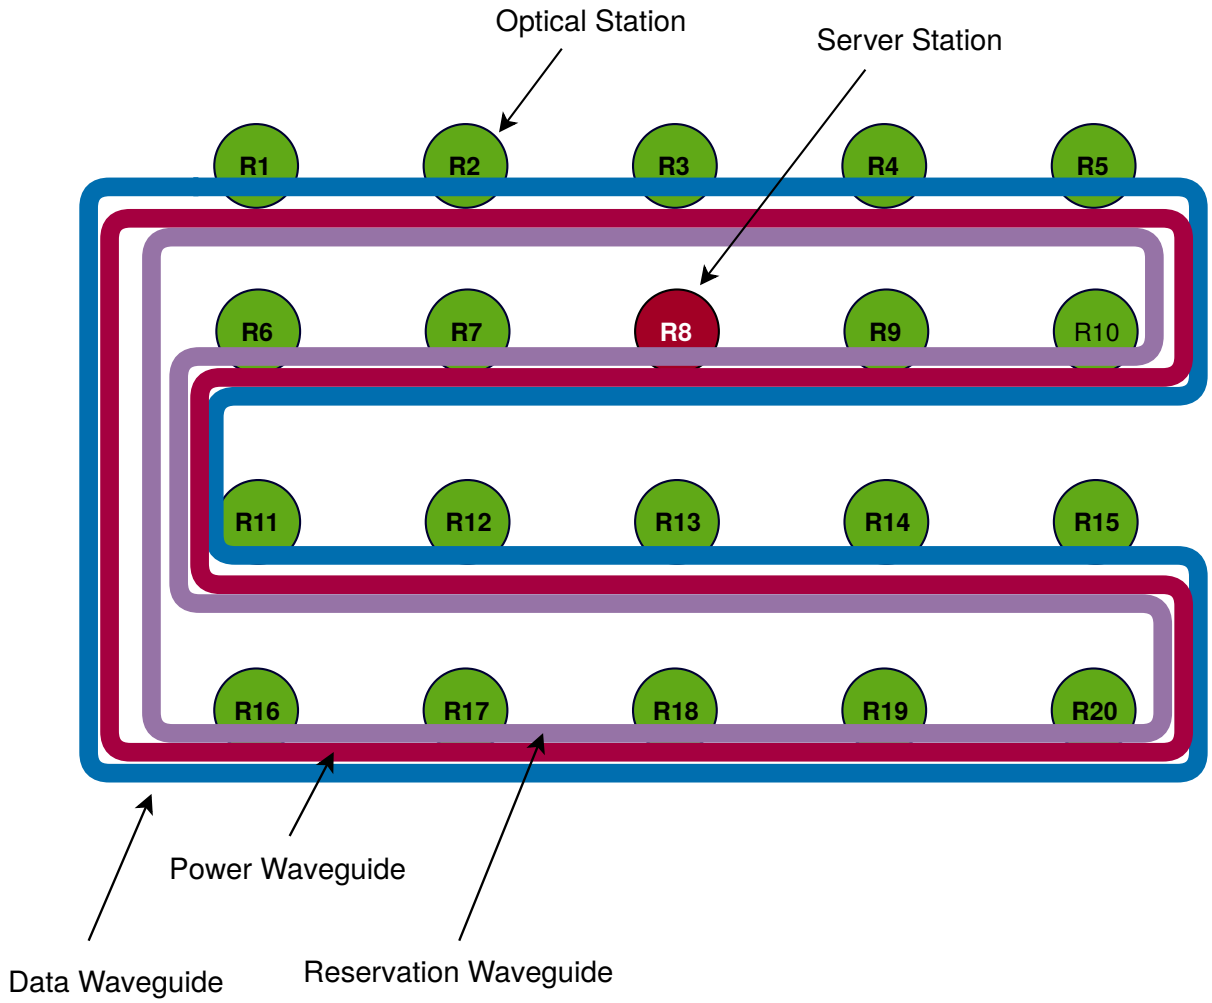

Figure 7: Layout of a serpentine shaped topology with data, power, and reservation waveguides

6.1 Communication framework between the representative server and an optical station

In our scheme, we propose that in order to send data through a data waveguide, the optical station first requests its respective representative server for the data waveguide. When it receives acknowledgement from its representative server then only it transfers data on the allotted data waveguide. The optical stations communicate with their respective representative servers through the reservation waveguides. Each optical station and its respective representative server is allotted some specific wavelengths in the reservation waveguide in order to communicate with each other. This set of wavelengths allotted to this pair is called its reservation channel. Each such pair in the network is allotted a separate reservation channel, thus, resulting in a logical point to point links between the optical stations and their respective representative servers. Whenever an optical station is required to communicate with its representative server, it writes its data on the respective reservation channel. The *Server Station* reads the data from the channel and accordingly forwards the data to the respective representative server. In our network, we require a maximum of 20 such reservation channels with each channel composed of 7 wavelengths. Thus, considering 64 wavelength waveguides, we require three such waveguides.

Each optical station and the *Server Station* have a set of in-situ on-chip Ge based lasers. These lasers are used to generate optical power whenever they want to communicate for reservation purpose. Moreover, we propose that the RSs are connected to each other through an electrical network. The main reasons behind using electrical network: close proximity of RSs and small size of messages.

Actions	Time Slots
Send request to RS	1
RSs make a decision	6
Parallel Activity	
Acknowledgement to the sender	1
Reservation message to the receiver	1
New Activity	
Data Transmission(64 bytes)	8
Total	16
1 Cycle = 2 Time slots	

Table 9: Operations for message transmission

6.2 Arbitration Messages

Whenever a station wants to send data, it first sends a message to its Representative Server (RS) at the center of the die. The station waits for the response from its RS before using any power or data waveguide. Table 9 shows the various operations and the number of cycles required by the optical station to send a full cache line (64 bytes). It should be noted that *PShaRe* uses a double pumping strategy like Corona [92] in which a station is allowed to transfer data at both the edges of the clock. The ability to send two messages in a single cycle is possible due to the use of ring resonators, which can potentially switch at a very high speed (>10GHz).

6.3 Representative Server Implementation

Each representative server is equivalent to a two port router connected to the representative cluster optical station using through silicon vias (TSV). In addition, the router uses an additional memory to store waveguide reservation table. Using the scaling rules proposed in [85], each such server acquires less than $0.7mm^2$ area (with $0.01mm^2$ required to store the waveguide reservation table).

7 Thermal tuning, Tuning Controller and Thermal Analysis

7.1 Thermal Tuning and Power Reuse

In photonic on-chip networks, microring resonators are the basic building blocks [4, 74]. They are used to insert and divert modulated wavelengths from the optical waveguides. Each microring resonator is sensitive to a particular wavelength, called *resonant wavelength*. Each resonator is designed to divert that particular wavelength only. However, the microring resonators are highly sensitive to temperature variations. Variation in the ring temperature results in the drift in the resonant wavelength (0.09nm/$^\circ C$) [4, 68]. Thus, large temperature fluctuations inside the chip can make the ring resonator to read a different wavelength, resulting in either the erroneous transmissions or in worst case the complete breakdown of the system. The standard approach used in prior work is to design all the ring resonators to operate at some maximum temperature τ and during the operational mode bring all the ring resonators to this maximum temperature. The temperature τ is chosen to be the maximum possible temperature that the chip may achieve.

To achieve this objective, the microring resonators are equipped with micro-heaters. These heaters are used to raise the temperature of the ring resonators to this maximum temperature τ so that they can work normally. This entire process of stabilizing the microring resonators under temperature variations is called tuning and the amount of power required is called *tuning power*.

In our scheme, we have tried to reduce the tuning power by reusing the unused optical power. However, it should be noted that we are not replacing the traditional approach of using the microring heaters. We are using both the schemes. Initially, we start to tune the ring resonators using the heat generated by photonic heaters by absorbing the unused optical power. However, if the heat generated is not enough to raise the temperature of the ring resonators

to the operating temperature τ, we take the help of microring heaters and generate the remaining heat. Thus, both the photonic heaters and micro-heaters are used to tune the ring resonators. Moreover, unlike traditional scheme of using a single micro-heater for a single ring resonator, we propose to use a single photonic heater for a group of co-located ring resonators. It is because our thermal simulation results showed that the co-located rings have a similar thermal profile and hence drift by the same amount. The same insight has also been derived in [68].

We want to mention here that to overcome the optical non-linearity in silicon waveguides, we are not using a single power or tuning waveguide to carry the entire power inside the chip. Instead the optical power generated by off-chip lasers is divided among 16 power waveguides, and at the end the unused power from these power waveguides is coupled into a set of tuning waveguides. Figure 8 shows such coupling.

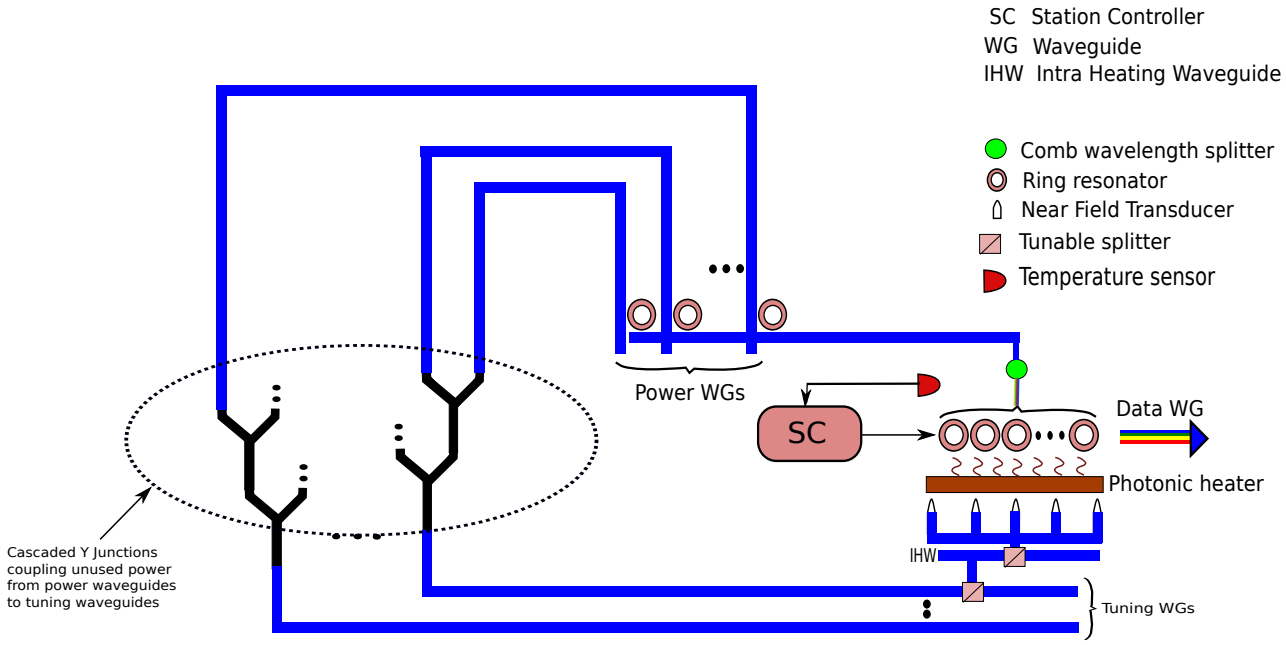

Figure 8: Power Reuse Scheme

7.2 Tuning Controller

Proportional-integral-derivative (PID) controllers are extremely popular and very well studied. They have been widely used in real world problems. A PID controller is made up of three controllers: **1) Proportional controller** which gives an output proportional to the error ($e(t)$), **2) Integral controller** which reduces the steady-state error and a **3) Derivative controller** that improves the transient response. The standard equation of a PID controller in the parallel sum form is given in Equation 6.

$$u(t) = K_p e(t) + K_i \int_0^t e(\tau)d\tau + K_d \frac{de(t)}{dt}, \tag{6}$$

where $u(t)$ is the control signal, $e(t)$ is the control error ($e = signal\ under\ observation - reference\ point$), K_p is the proportional gain, K_i is the integral gain and K_d is the derivative gain. In the discrete domain Equation 6 changes to Equation 7.

$$U(z) = K_p E(z) + \frac{K_i}{1-z^{-1}} + K_d(1-z^{-1})E(z) \tag{7}$$

In the difference form Equation 7 can be written as:

$$u(n) = u(n-1) + K_p(e(n) - e(n-1)) + K_i T_s e(n) + \frac{K_d}{T_s}(e(n) - 2e(n-1) + e(n-2)), \tag{8}$$

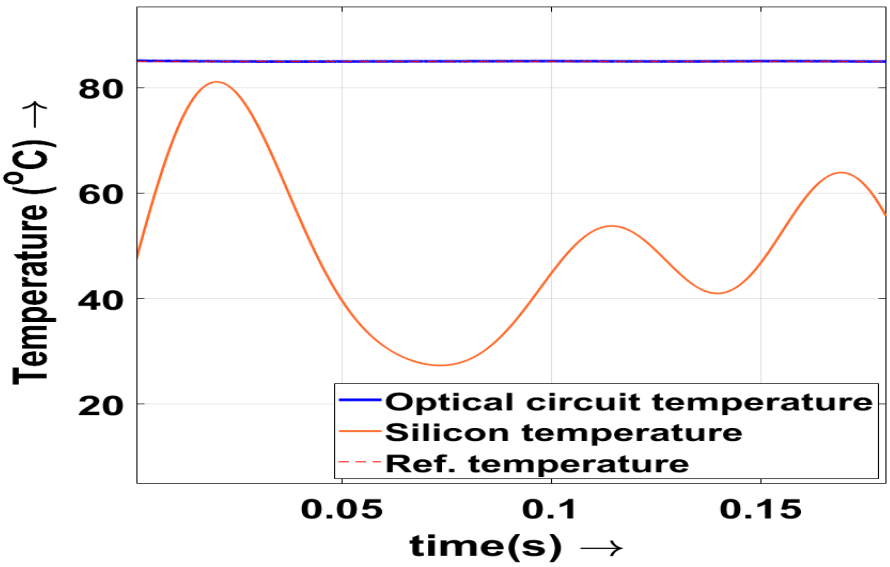

Figure 9: Temperatures in the silicon and optical layers

where T_s is the sampling period and n is the index. To model the thermal response of the chip, we made a model in COMSOL5.3a (CFD solver) and extracted the equation for the transient response, which is given by Equation 9 for a $1W$ circular source of radius $100\mu m$.

$$\Delta T = 19.5 - 2.186e^{\frac{-t}{2.29\times10^{-3}}} - 13.26e^{\frac{-t}{5.18\times10^{-5}}} \tag{9}$$

We use this model as the plant to tune the parameters of the PID controller in Simulink. The parameters obtained form Simulink are shown in Table 10.

K_p	877.79
K_i	36540418.36
K_d	-0.01146

Table 10: PID controller parameters

We implemented Equation 8 in VHDL and synthesized it using the Cadence RTL compiler for the 14 nm technology node. The area and frequency obtained were $\approx 4000\mu m^2$ and $1.35GHz$ respectively. Figure 9 shows the silicon and optical circuits' temperature. The temperature remains roughly constant (error of 0.1°C).

7.3 Thermal Analysis

To model the thermal variations inside the chip, we use the thermal simulator, HotSpot [103]. It takes the floorplan of the chip and the power consumption (power trace) of all the blocks (functional units, cores, and caches) inside the chip as input and yields the temperature profile.

The floorplan of our chip used for thermal modelling in HotSpot is shown in Figure 10. These are the basic blocks inside the chip:

1. CPU cores
2. Cache banks
3. Optical stations
4. Memory controllers

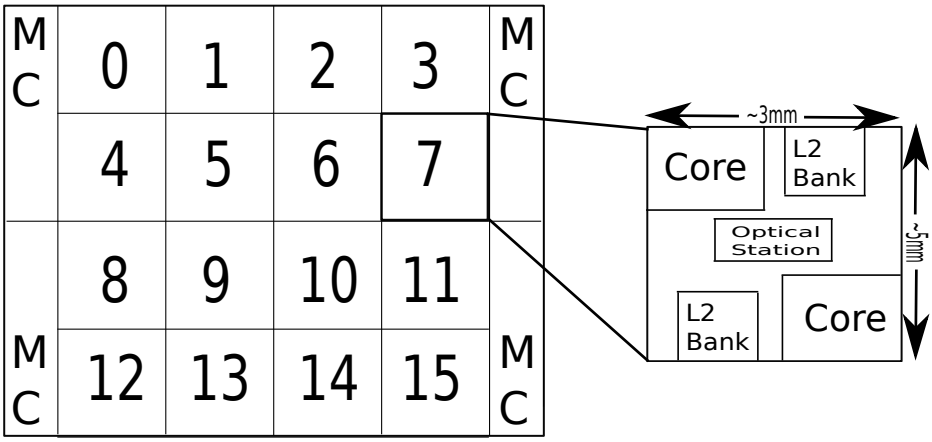

Figure 10: (Left)Floorplan for HotSpot. (Right)One Tile in a floorplan

To simulate the power and energy consumption in our network, we use the Tejas simulator, which is bundled with the Orion and McPAT tools. Tejas provides the laser activity and based on this activity we analytically calculated the power consumption as described in Section 5.2.2.

We collect the power values from the Tejas simulator for each benchmark. These values and the floorplan of the chip are provided as inputs to HotSpot [103] that calculates the temperature variations over the chip. The temperature variations change the leakage power in the chip. As a result, the procedure is repeated iteratively until the system reaches a stable state.

PShaRe has a photonic layer underneath the logic layer. The photonic layer carries the optical signals and is associated with optical power losses. The chip temperature depends on the amount of power consumed by cores, cache banks and other logic layer components, and the optical power consumption. Tejas provides all of these power values – both electrical and optical.

Table 11 shows the temperature variations across the chip for different benchmarks and the amount of optical power that remains unused inside the chip. We observe that on an average we have 2.9W of optical power that remains unused inside the chip and hence can be used to heat the ring resonators.

Benchmark	**Temperature Range($^{\circ}C$)**	**Unused Power (W)**	**Reduction in Trimming power (%)**
barnes	45-71	3.6	43
blackscholes	45-65	2.6	34
bodytrack	45-63	2.7	37
canneal	45-61	3.1	38
cholesky	45-73	3.9	49
dedup	45-56	2.3	33
ferret	45-60	2.45	34
fft	45-76	3.7	44
fluidanimate	45-70	2.6	38
fmm	45-72	3.8	45
ocean_ncp	45-75	3.3	41
radiosity	45-70	2.3	29
streamcluster	45-60	2.34	32
Mean	45-68	2.9	38.2

Table 11: Thermal simulations and unused power

References

[1] acacia. Acacia communications. Https://acacia-inc.com/.

[2] A. Ardakani, F. Leduc-Primeau, N. Onizawa, T. Hanyu, and W. J. Gross, "Vlsi implementation of deep neural network using integral stochastic computing," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, Oct 2017.

[3] S. Assefa, S. Shank, W. Green, M. Khater, E. Kiewra, C. Reinholm, S. Kamlapurkar, A. Rylyakov, C. Schow, F. Horst, H. Pan, T. Topuria, P. Rice, D. M. Gill, J. Rosenberg, T. Barwicz, M. Yang, J. Proesel, J. Hofrichter, B. Offrein, X. Gu, W. Haensch, J. Ellis-Monaghan, and Y. Vlasov, "A 90nm cmos integrated nano-photonics technology for 25gbps wdm optical communications applications," in *2012 International Electron Devices Meeting*, Dec 2012.

[4] J. Bashir, Eldhose, and S. R. Sarangi, "A survey of on-chip optical interconnects," *ACM Comput. Surv.*, vol. 51, no. 6, jan 2019.

[5] J. Bashir, E. Peter, and S. R. Sarangi, "Bigbus: A scalable optical interconnect," *J. Emerg. Technol. Comput. Syst.*, vol. 15, no. 1, jan 2019.

[6] J. Bashir and S. R. Sarangi, "Nuplet: A photonics based multi-chip nuca architecture," in *2017 IEEE 35th International Conference on Computer Design (ICCD)*. IEEE, 2017.

[7] C. Batten, A. Joshi, J. Orcutt, A. Khilo, B. Moss, C. Holzwarth, M. Popovic, H. Li, H. I. Smith, J. Hoyt *et al.*, "Building manycore processor-to-dram networks with monolithic silicon photonics," in *High Performance Interconnects (HOTI)*. IEEE, 2008, pp. 21–30.

[8] A. Bianco, D. Cuda, M. Garrich, G. G. Castillo, R. Gaudino, and P. Giaccone, "Optical interconnection networks based on microring resonators," *Optical Communications and Networking, IEEE/OSA Journal of*, vol. 4, no. 7, pp. 546–556, 2012.

[9] A. Biberman, K. Preston, G. Hendry, N. Sherwood-Droz, J. Chan, J. S. Levy, M. Lipson, and K. Bergman, "Photonic network-on-chip architectures using multilayer deposited silicon materials for high-performance chip multiprocessors," *ACM Journal on Emerging Technologies in Computing Systems (JETC)*, vol. 7, no. 2, p. 7, 2011.

[10] C. Bienia, S. Kumar, J. P. Singh, and K. Li, "The PARSEC benchmark suite: characterization and architectural implications," in *PACT*, 2008.

[11] J.-R. Burie, G. Beuchet, M. Mimoun, P. Pagnod-Rossiaux, B. Ligat, J. Bertreux, J.-M. Rousselet, J. Dufour, P. Rougeolle, and F. Laruelle, "Ultra high power, ultra low rin up to 20 ghz 1.55 μm dfb algainasp laser for analog applications," in *OPTO*. International Society for Optics and Photonics, 2010, pp. 76 160Y–76 160Y.

[12] R. E. Camacho-Aguilera, Y. Cai, N. Patel, J. T. Bessette, M. Romagnoli, L. C. Kimerling, and J. Michel, "An electrically pumped germanium laser," *Optics express*, vol. 20, no. 10, pp. 11 316–11 320, 2012.

[13] J. Cardenas, C. B. Poitras, J. T. Robinson, K. Preston, L. Chen, and M. Lipson, "Low loss etchless silicon photonic waveguides," *Opt. Express*, Mar 2009.

[14] W. A. Challener, E. Gage, A. Itagi, and C. Peng, "Optical transducers for near field recording," *Japanese journal of applied physics*, 2006.

[15] H.-H. Chang, A. W. Fang, M. N. Sysak, H. Park, R. Jones, O. Cohen, O. Raday, M. J. Paniccia, and J. E. Bowers, "1310nm silicon evanescent laser," *Optics Express*, vol. 15, no. 18, pp. 11 466–11 471, 2007.

[16] K.-N. Chen, M. J. Kobrinsky, B. C. Barnett, and R. Reif, "Comparisons of conventional, 3-d, optical, and rf interconnects for on-chip clock distribution," *Electron Devices, IEEE Transactions on*, vol. 51, no. 2, pp. 233–239, 2004.

[17] S. V. R. Chittamuru, S. Desai, and S. Pasricha, "Swiftnoc: A reconfigurable silicon-photonic network with multicast-enabled channel sharing for multicore architectures," *J. Emerg. Technol. Comput. Syst.*, Jun. 2017.

[18] M. J. Cianchetti, J. C. Kerekes, and D. H. Albonesi, "Phastlane: a rapid transit optical routing network," in *ACM SIGARCH Computer Architecture News*, vol. 37, no. 3. ACM, 2009, pp. 441–450.

[19] Z. J. Coppens, W. Li, D. G. Walker, and J. G. Valentine, "Probing and controlling photothermal heat generation in plasmonic nanostructures," *Nano letters*, vol. 13, no. 3, 2013.

[20] R. Das, S. Eachempati, A. K. Mishra, V. Narayanan, and C. R. Das, "Design and evaluation of a hierarchical on-chip interconnect for next-generation cmps," in *2009 IEEE 15th International Symposium on High Performance Computer Architecture*, Feb 2009, pp. 175–186.

[21] Y. Demir and N. Hardavellas, "Slac: Stage laser control for a flattened butterfly network," in *2016 IEEE International Symposium on High Performance Computer Architecture (HPCA)*. IEEE, 2016, pp. 321–332.

[22] Y. Demir, Y. Pan, S. Song, N. Hardavellas, J. Kim, and G. Memik, "Galaxy: A high-performance energy-efficient multi-chip architecture using photonic interconnects," in *Proceedings of the 28th ACM international conference on Supercomputing*. ACM, 2014, pp. 303–312.

[23] C. T. DeRose, M. R. Watts, D. C. Trotter, D. L. Luck, G. N. Nielson, and R. W. Young, "Silicon microring modulator with integrated heater and temperature sensor for thermal control." Optical Society of America, 2010.

[24] A. W. Fang, B. R. Koch, R. Jones, E. Lively, D. Liang, Y.-H. Kuo, and J. E. Bowers, "A distributed bragg reflector silicon evanescent laser," *IEEE Photonics Technology Letters*, vol. 20, no. 20, pp. 1667–1669, 2008.

[25] A. W. Fang, H. Park, O. Cohen, R. Jones, M. J. Paniccia, and J. E. Bowers, "Electrically pumped hybrid algainas-silicon evanescent laser," *Optics express*, vol. 14, no. 20, pp. 9203–9210, 2006.

[26] Q. Fang, J. Song, X. Luo, L. Jia, M. Yu, G. Lo, and Y. Liu, "High efficiency ring-resonator filter with nisi heater," *IEEE Photonics Technology Letters*, March 2012.

[27] M. FAUGERON, M. Chtioui, A. Enard, O. Parillaud, F. Lelarge, M. Achouche, J. Jacquet, A. Marceaux, and F. van Dijk, "High optical power, high gain and high dynamic range directly modulated optical link," *Lightwave Technology, Journal of*, vol. 31, no. 8, pp. 1227–1233, April 2013.

[28] M. Faugeron, M. Tran, O. Parillaud, M. Chtioui, Y. Robert, E. Vinet, A. Enard, J. Jacquet, and F. V. Dijk, "High-power tunable dilute mode dfb laser with low rin and narrow linewidth," *Photonics Technology Letters, IEEE*, vol. 25, no. 1, pp. 7–10, 2013.

[29] M. Faugeron, M. Tran, F. Lelarge, M. Chtioui, Y. Robert, E. Vinet, A. Enard, J. Jacquet, and F. Van Dijk, "High-power, low rin 1.55-directly modulated dfb lasers for analog signal transmission," *Photonics Technology Letters*, vol. 24, no. 2, pp. 116–118, 2012.

[30] G. Fischbeck, R. Moosburger, C. Kostrzewa, A. Achen, and K. Petermann, "Singlemode optical waveguides using a high temperature stable polymer with low losses in the 1.55 mu;m range," *Electronics Letters*, Mar 1997.

[31] T. Fukamachi, K. Adachi, K. Shinoda, S. Tsuji, T. Kitatani, S. Tanaka, and M. Aoki, "Recent progress in 1.3-μm uncooled ingaalas directly modulated lasers," in *Semiconductor Laser Conference (ISLC), 2010 22nd IEEE International*. IEEE, 2010, pp. 189–190.

[32] K. M. Geib, K. D. Choquette, A. Allerman, J. J. Hindi, J. Nevers, and B. E. Hammous, "Monolithically integrated vcsels and photodetectors for microsystem applications," in *Lasers and Electro-Optics Society Annual Meeting, 1998. LEOS '98. IEEE*, vol. 2, Dec 1998, pp. 27–28 vol.2.

[33] R. R. Ghosh, J. Bashir, S. R. Sarangi, and A. Dhawan, "Spliesr: Tunable power splitter based on an electro-optic slotted ring resonator," *Optics Communications*, 2019.

[34] E. Griese, "A high-performance hybrid electrical-optical interconnection technology for high-speed electronic systems," *Advanced Packaging, IEEE Transactions on*, vol. 24, no. 3, pp. 375–383, 2001.

[35] H. Gu and J. Xu, "Design of 3d optical network on chip," in *2009 Symposium on Photonics and Optoelectronics*, Aug 2009.

[36] 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)*, Sept 2014.

[37] J. Hao, L. Zhou, and M. Qiu, "Nearly total absorption of light and heat generation by plasmonic metamaterials," *Phys. Rev. B*, Apr 2011.

[38] M. Haurylau, G. Chen, H. Chen, J. Zhang, N. A. Nelson, D. H. Albonesi, E. G. Friedman, and P. M. Fauchet, "On-chip optical interconnect roadmap: challenges and critical directions," *Selected Topics in Quantum Electronics, IEEE Journal of*, vol. 12, no. 6, pp. 1699–1705, 2006.

[39] M. Horowitz, "1.1 computing's energy problem (and what we can do about it)," in *2014 IEEE International Solid-State Circuits Conference Digest of Technical Papers (ISSCC)*, Feb 2014.

[40] HP. (2008) Rings of fire. Http://www.hpl.hp.com/news/2008/oct-dec/photonics2.html.

[41] J.-S. Huang, H. Lu, and H. Su, "Ultra-high power, low rin and narrow linewidth lasers for 1550nm dwdm 100km long-haul fiber optic link," in *IEEE Lasers and Electro-Optics Society, 2008. LEOS 2008. 21st Annual Meeting of the*, Nov 2008, pp. 894–895.

[42] A. Joshi, C. Batten, Y.-J. Kwon, S. Beamer, I. Shamim, K. Asanovic, and V. Stojanovic, "Silicon-photonic clos networks for global on-chip communication," in *NoCS*, 2009.

[43] P. Kapur and K. C. Saraswat, "Comparisons between electrical and optical interconnects for on-chip signaling," in *Interconnect Technology Conference, 2002. Proceedings of the IEEE 2002 International*. IEEE, 2002, pp. 89–91.

[44] T. J. Kippenberg, J. Kalkman, A. Polman, and K. J. Vahala, "Demonstration of an erbium-doped microdisk laser on a silicon chip," *Physical Review A*, vol. 74, no. 5, p. 051802, 2006.

[45] N. Kirman and J. F. Martínez, "A power-efficient all-optical on-chip interconnect using wavelength-based oblivious routing," in *ACM Sigplan Notices*, vol. 45, no. 3. ACM, 2010, pp. 15–28.

[46] M. J. Kobrinsky, "On-chip optical interconnects," *Intel Technology Journal*, vol. 8, no. 2, pp. 129–142, 2004.

[47] P. Koka, M. O. McCracken, H. Schwetman, C.-H. O. Chen, X. Zheng, R. Ho, K. Raj, and A. V. Krishnamoorthy, "A micro-architectural analysis of switched photonic multi-chip interconnects," in *ACM SIGARCH Computer Architecture News*, vol. 40, no. 3. IEEE Computer Society, 2012, pp. 153–164.

[48] P. Koka, M. O. McCracken, H. Schwetman, X. Zheng, R. Ho, and A. V. Krishnamoorthy, "Silicon-photonic network architectures for scalable, power-efficient multi-chip systems," in *ACM SIGARCH Computer Architecture News*, vol. 38, no. 3. ACM, 2010, pp. 117–128.

[49] S. Koohi and S. Hessabi, "All-optical wavelength-routed architecture for a power-efficient network on chip," *Computers, IEEE Transactions on*, vol. 63, no. 3, pp. 777–792, 2014.

[50] E. Kotelnikov, A. Katsnelson, K. Patel, and I. Kudryashov, "High-power single-mode ingaasp/inp laser diodes for pulsed operation," in *Novel In-Plane Semiconductor Lasers XI*, vol. 8277. International Society for Optics and Photonics, 2012, p. 827715.

[51] G. Kurian, J. E. Miller, J. Psota, J. Eastep, J. Liu, J. Michel, L. C. Kimerling, and A. Agarwal, "Atac: a 1000-core cache-coherent processor with on-chip optical network," in *PACT*, 2010.

[52] J. A. Lazaro, J. Gonzalez, J. A. Altabas, and A. Lerin, "Graphene silicon ring resonators for wavelength routers in photonic network-on-chip," in *2015 17th International Conference on Transparent Optical Networks (ICTON)*, July 2015.

[53] J. S. Levy, Y. Okawachi, M. Lipson, A. L. Gaeta, and K. Saha, "High-performance silicon-based multiple wavelength source," in *CLEO: Science and Innovations*. OSA, 2011, p. CMAA7.

[54] C. Li, C.-H. Chen, B. Wang, S. Palermo, M. Fiorentino, and R. Beausoleil, "Design of an energy-efficient silicon microring resonator-based photonic transmitter," *no. August*, pp. 46–54, 2014.

[55] C. Li, P. V. Gratz, and S. Palermo, "Nano-photonic networks-on-chip for future chip multiprocessors," in *More than Moore Technologies for Next Generation Computer Design*. Springer, 2015, pp. 155–186.

[56] Z. Li and T. Li, "Espn: A case for energy-star photonic on-chip network," in *Proceedings of the 2013 International Symposium on Low Power Electronics and Design*, ser. ISLPED '13. IEEE Press, 2013, pp. 377–382.

[57] E. Lippert, H. Fonnum, G. Arisholm, and K. Stenersen, "A 22-watt mid-infrared optical parametric oscillator with v-shaped 3-mirror ring resonator," *Optics express*, 2010.

[58] J. Liu, X. Sun, R. Camacho-Aguilera, L. C. Kimerling, and J. Michel, "Ge-on-si laser operating at room temperature," *Optics letters*, vol. 35, no. 5, pp. 679–681, 2010.

[59] Lumerical, "Lumerical mode sol." [Online]. Available: https://https://www.lumerical.com/tcadproducts/mode/

[60] LUXTERA. Luxtera: Fibre to the chip. Http://www.luxtera.com/luxtera/products.

[61] P. K. Meher, "An optimized lookup-table for the evaluation of sigmoid function for artificial neural networks," in *2010 18th IEEE/IFIP International Conference on VLSI and System-on-Chip*, Sept 2010.

[62] Mellanox. Mellanox technologies. Http://www.mellanox.com/.

[63] R. Michalzik, *VCSELs: fundamentals, technology and applications of vertical-cavity surface-emitting lasers*. Springer, 2012, vol. 166.

[64] K. H. Mo, Y. Ye, X. Wu, W. Zhang, W. Liu, and J. Xu, "A hierarchical hybrid optical-electronic network-on-chip," in *VLSI (ISVLSI), 2010 IEEE Computer Society Annual Symposium on*. IEEE, 2010, pp. 327–332.

[65] R. Morris, E. Jolley, and A. K. Kodi, "Extending the performance and energy-efficiency of shared memory multicores with nanophotonic technology," *Parallel and Distributed Systems, IEEE Transactions on*, vol. 25, no. 1, pp. 83–92, 2014.

[66] N. Muralimanohar, R. Balasubramonian, and N. P. Jouppi, "Cacti 6.0: A tool to model large caches," *HP laboratories*, pp. 22–31, 2009.

[67] B. News. (2015) Engineers demo first processor that uses light for ultrafast communications. Http://news.berkeley.edu/2015/12/23/electronic-photonic-microprocessor-chip/.

[68] C. Nitta, M. Farrens, and V. Akella, "Addressing system-level trimming issues in on-chip nanophotonic networks," in *HPCA*, Feb 2011, pp. 122–131.

[69] Y. Pan, J. Kim, and G. Memik, "Flexishare: Channel sharing for an energy-efficient nanophotonic crossbar," in *HPCA*, 2010.

[70] Y. Pan, P. Kumar, J. Kim, G. Memik, Y. Zhang, and A. Choudhary, "Firefly: illuminating future network-on-chip with nanophotonics," in *ACM SIGARCH Computer Architecture News*. ACM, 2009.

[71] M. Paniccia and J. Bowers, "First electrically pumped hybrid pumped hybrid silicon laser silicon laser," 2006, http://www.intel.com/content/dam/www/public/us/en/documents/technology-briefs/intel-labs-hybridsilicon-laser-announcement.pdf.

[72] S. Pasricha and S. Bahirat, "Opal: A multi-layer hybrid photonic noc for 3d ics," in *Design Automation Conference (ASP-DAC), 2011 16th Asia and South Pacific*. IEEE, 2011, pp. 345–350.

[73] E. Peter, A. Arora, J. Bashir, A. Bagaria, and S. R. Sarangi, "Optical overlay nuca: A high-speed substrate for shared l2 caches," *J. Emerg. Technol. Comput. Syst.*, vol. 13, no. 4, 2017.

[74] E. Peter, A. Thomas, A. Dhawan, and S. R. Sarangi, "Coldbus: A near-optimal power efficient optical bus," in *HiPC*, 2015.

[75] K. Petermann, *Laser diode modulation and noise*. Springer Science & Business Media, 2012, vol. 3.

[76] J. Puche, S. Lechago, S. Petit, M. E. Gómez, and J. Sahuquillo, "Accurately modeling a photonic noc in a detailed cmp simulation framework," in *High Performance Computing & Simulation (HPCS), 2016 International Conference on*. IEEE, 2016, pp. 387–394.

[77] X. Qianfan, M. Sasikanth, S. Brad, S. Jagat, and L. Michal, "12.5 gbit/s carrier-injection-based silicon micro-ring silicon modulators," *Opt. Express*, vol. 15, Jan 2007.

[78] J. Ralston, D. Gallagher, P. Tasker, H. Zappe, I. Esquivias, and J. Fleissner, "Vertically compact 15 ghz gaas/algaas multiple quantum well laser grown by molecular beam epitaxy," *Electronics Letters*, vol. 27, no. 19, pp. 1720–1722, 1991.

[79] G. T. Reed, *Silicon Photonics: The State of the Art*. John Wiley & Sons, 2008.

[80] H. Rong, S. Xu, O. Cohen, O. Raday, M. Lee, V. Sih, and M. Paniccia, "A cascaded silicon raman laser," *Nature photonics*, vol. 2, no. 3, p. 170, 2008.

[81] A. Sakai, T. Fukazawa, and T. Baba, "Low loss ultra-small branches in a silicon photonic wire waveguide," *IEICE Transactions on Electronics*, 2002.

[82] S. R. Sarangi, K. Rajshekar, K. Prathmesh, G. Seep, and P. Eldhose, "Tejas: A java based versatile micro-architectural simulator,," in *PATMOS*, 2015.

[83] sicoya. Silicon photonics chips. Http://sicoya.com/.

[84] L. sol, "Lumerical device." [Online]. Available: https://https://www.lumerical.com/tcadproducts/device/

[85] M. R. Stan, K. Skadron, W. Huang, and K. Rajamani, "Scaling with design constraints: Predicting the future of big chips," *IEEE Micro*, vol. 31, 2011.

[86] S. Stankovic, R. Jones, M. N. Sysak, J. M. Heck, G. Roelkens, and D. Van Thourhout, "1310-nm hybrid iii–v/si fabry–pérot laser based on adhesive bonding," *IEEE Photonics Technology Letters*, vol. 23, no. 23, pp. 1781–1783, 2011.

[87] SYNOPSYS, "Synopsys rsoft," https://www.synopsys.com/optical-solutions/rsoft.html.

[88] M. Tan, P. Rosenberg, J. S. Yeo, M. McLaren, S. Mathai, T. Morris, H. P. Kuo, J. Straznicky, N. P. Jouppi, and S.-Y. Wang, "A high-speed optical multi-drop bus for computer interconnections," *Applied Physics A*, vol. 95, no. 4, pp. 945–953, 2009.

[89] E. TECH. (2014) Hp bets it all on the machine, a new computer architecture based on memristors and silicon photonics. Http://www.extremetech.com/extreme/184165-hp-bets-it-all-on-the-machine-a-new-computer-architecture-based-on-memristors-and-silicon-photonics/.

[90] A. W. Topol, D. C. L. Tulipe, L. Shi, D. J. Frank, K. Bernstein, S. E. Steen, A. Kumar, G. U. Singco, A. M. Young, K. W. Guarini, and M. Ieong, "Three-dimensional integrated circuits," *IBM Journal of Research and Development*, July 2006.

[91] E. Vahapoglu and M. Altun, "Accurate synthesis of arithmetic operations with stochastic logic," in *2016 IEEE Computer Society Annual Symposium on VLSI (ISVLSI)*, July 2016.

[92] D. Vantrease, R. Schreiber, M. Monchiero, M. McLaren, N. P. Jouppi, M. Fiorentino, A. Davis, N. Binkert, R. G. Beausoleil, and J. H. Ahn, "Corona: System Implications of Emerging Nanophotonic Technology," in *ISCA*, 2008.

[93] L. Vivien, F. Grillot, E. Cassan, D. Pascal, S. Lardenois, A. Lupu, S. Laval, M. Heitzmann, and J.-M. Fdli, "Comparison between strip and rib {SOI} microwaveguides for intra-chip light distribution," *Optical Materials*, vol. 27, no. 5, pp. 756 – 762, 2005, si-based Photonics: Towards True Monolithic IntegrationProceedings of the European Materials Research Society Symposium {A1European} Materials Research Society 2004 Spring Meeting. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S092534670400223X

[94] S. Werner, J. Navaridas, and M. Lujn, "Amon: An advanced mesh-like optical noc," in *2015 IEEE 23rd Annual Symposium on High-Performance Interconnects*, Aug 2015.

[95] S. V. 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)*, Feb 2018.

[96] P. Wolf, P. Moser, G. Larisch, W. Hofmann, H. Li, J. A. Lott, C.-Y. Lu, S. L. Chuang, and D. Bimberg, "Energy-efficient and temperature-stable high-speed vcsels for optical interconnects," in *Transparent Optical Networks (ICTON), 2013 15th International Conference on*. IEEE, 2013, pp. 1–5.

[97] S. C. Woo, M. Ohara, E. Torrie, J. P. Singh, and A. Gupta, "The splash-2 programs: characterization and methodological considerations," *SIGARCH Comput. Archit. News*, vol. 23, pp. 24–36, May 1995.

[98] X. Wu, J. Xu, Y. Ye, X. Wang, M. Nikdast, Z. Wang, and Z. Wang, "An inter/intra-chip optical network for manycore processors," *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, 2015.

[99] Q. Xu, B. Schmidt, S. Pradhan, and M. Lipson, "Micrometre-scale silicon electro-optic modulator," *Nature*, May 2005.

[100] L. Yang, T. Carmon, B. Min, S. M. Spillane, and K. J. Vahala, "Erbium-doped and raman microlasers on a silicon chip fabricated by the sol–gel process," *Applied Physics Letters*, vol. 86, no. 9, p. 091114, 2005.

[101] Y. Ye, L. Duan, J. Xu, J. Ouyang, M. K. Hung, and Y. Xie, "3d optical networks-on-chip (noc) for multiprocessor systems-on-chip (mpsoc)," in *2009 IEEE International Conference on 3D System Integration*, Sept 2009.

[102] Y. Ye, J. Xu, B. Huang, X. Wu, W. Zhang, X. Wang, M. Nikdast, Z. Wang, W. Liu, and Z. Wang, "3-d mesh-based optical network-on-chip for multiprocessor system-on-chip," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, April 2013.

[103] R. Zhang, M. R. Stan, and K. Skadron, "Hotspot 6.0: Validation, acceleration and extension," 2015.

[104] L. Zhou and A. K. Kodi, "Probe: Prediction-based optical bandwidth scaling for energy-efficient nocs," in *NOCS*, 2013.

[105] A. Zulfiqar, P. Koka, H. Schwetman, M. Lipasti, X. Zheng, and A. Krishnamoorthy, "Wavelength stealing: an opportunistic approach to channel sharing in multi-chip photonic interconnects," in *MICRO*, 2013.