

# Sharing and Placement of On-chip Laser Sources in Silicon-Photonic NoCs

Chao Chen, Tiansheng Zhang, Pietro Contu, Jonathan Klamkin, Ayse K. Coskun, Ajay Joshi  
 Electrical and Computer Engineering Department, Boston University, Boston, MA, USA  
 {chao9810, tszhang, contu, klamkin, acoskun, joshi}@bu.edu

**Abstract**—Silicon-photonic links are projected to replace the electrical links for global on-chip communications in future manycore systems. The use of off-chip laser sources to drive these silicon-photonic links can lead to higher link losses, thermal mismatch between laser source and on-chip photonic devices, and packaging challenges. Therefore, on-chip laser sources are being evaluated as candidates to drive the on-chip photonic links. In this paper, we first explore the power, efficiency and temperature tradeoffs associated with an on-chip laser source. Using a 3D stacked system that integrates a manycore chip with the optical devices and laser sources, we explore the design space for laser source sharing (among waveguides) and placement to minimize laser power by simultaneously considering the network bandwidth requirements, thermal constraints, and physical layout constraints. As part of this exploration we consider Clos and crossbar logical topologies, U-shaped and W-shaped physical layouts, and various sharing/placement strategies: locally-placed dedicated laser sources for waveguides, locally-placed shared laser sources, and shared laser sources placed remotely along the chip edges. Our analysis shows that logical topology, physical layout, and photonic device losses strongly drive the laser source sharing and placement choices to minimize laser power.

## I. INTRODUCTION

Silicon-photonic link technology has been extensively explored as a potential replacement to electrical link technology in the design of network-on-chip (NoC) for manycore systems [1]–[6]. The key advantages of silicon-photonic link technology include up to an order of magnitude higher bandwidth density (compared to electrical link technology) through dense wavelength division multiplexing (DWDM), and link length independent data-dependent power. However, this silicon-photonic link technology has not yet been adopted for designing the NoC of commercial manycore systems as the fixed power consumed by the laser sources and the power consumed in thermal management of the silicon-photonic links can be significant and can negate the bandwidth density advantages. Moreover, packaging such a system where several off-chip laser sources drive the photonic NoC can be extremely challenging. Using on-chip laser sources to drive the photonic NoC is therefore being considered as a potential alternative [7], [8]. Although on-chip lasers require further technological development, they already provide advantages including the elimination of coupling losses and simplified packaging. On-chip lasers can also be switched on and off relatively rapidly and at a lower energy cost, which makes them compatible with various run-time laser power management techniques [9]–[11] that have been proposed. Lastly, given that laser sources and photonic component can exist in adjacent layers, it is possible to match the temperatures among these components, which would simplify thermal management techniques [12]–[15].

A current limitation of on-chip laser sources is the low wall-plug efficiency, which depends on both the optical output

power and temperature of the laser source. The optical power that needs to be output by a laser source depends on the physical layout of the silicon-photonic NoC and the bandwidth (i.e., number of required wavelengths) of the NoC channels. The laser source efficiency is low when the optical output power is too high or too low. Considering existing laser and silicon-photonic device technologies, the optical output power corresponding to the laser source's peak efficiency is likely higher than that required for each DWDM channel. Therefore, sharing of a laser source across parallel waveguides would be required to operate the laser source at its peak efficiency. At the same time, the temperature of each laser source is defined by the power consumed by the hardware in the laser source's neighborhood. As the laser temperature increases, the efficiency of the laser source decreases. Hence, the laser sources need to be strategically placed to operate them at as minimum a temperature as possible to maximize the laser source efficiency.

In this paper, we explore the design space of sharing and placement of on-chip laser sources by simultaneously considering the NoC bandwidth constraints driven by the applications running on the manycore system, thermal constraints driven by the power consumed by the cores and the laser sources, and the physical layout constraints driven by the losses in the photonic devices. The goal is to optimize the sharing as well as the placement of on-chip laser sources to maximize laser efficiency and in turn minimize the electrical power consumption of the lasers. Using a 256-core 3D-integrated system consisting of separate processor logic layer, photonic device layer, and laser source layer, we show that, for various NoC logical topologies and physical layouts, laser power consumption can be lowered by sharing laser sources across the various silicon-photonic waveguides and by intelligently placing these laser sources. It should be noted that our proposed approach is also valid for a system where the photonic devices are monolithically integrated with the CMOS devices.

The rest of the paper is organized as follows. Section II discusses the related work, followed by a description of the experimental setup that is used for our case study in Section III. In Section IV, we describe our methodology for optimizing the operating points for on-chip laser sources, and then we evaluate our methodology for various logical topologies and physical layouts of the silicon-photonic NoC in Section V.

## II. RELATED WORK

Several on-chip laser source technologies have been developed in recent years [16]–[19]. Although these technologies require further development, they show promise for simplifying packaging and enhancing the performance of photonic NoCs. On-chip laser sources can also be switched on and off

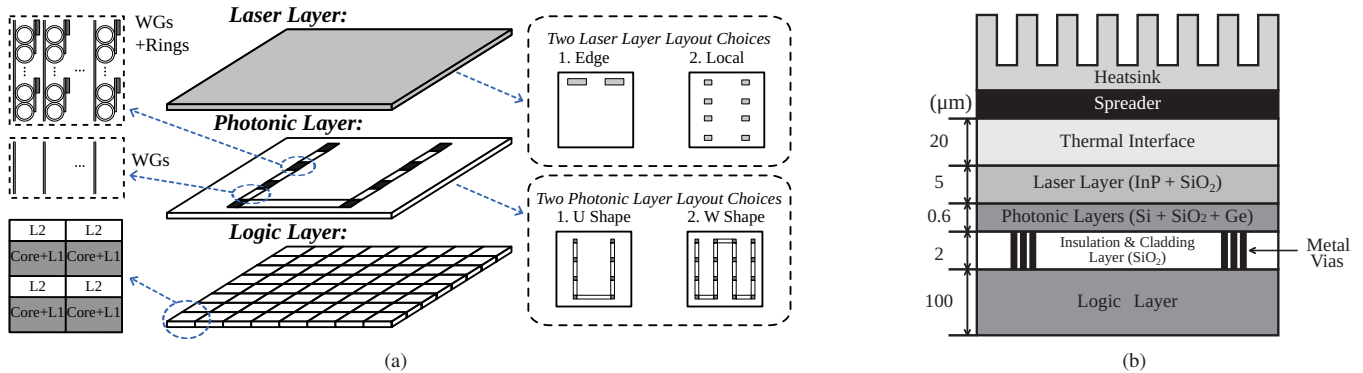


Fig. 1: (a) Overview of the 3D manycore system and layouts for each layer. Output from the laser sources is routed into the waveguides through couplers. The ring modulator is driven through TSVs by the transmitter in the logic layer. Photodetector's output is fed to the receiver on the logic layer through TSVs. (b) Cross-sectional view of our target 3D manycore system.

more efficiently compared to off-chip laser sources. This eases the process of run-time laser power management. In [7], the authors evaluate their ATAC photonic NoC architecture that is driven using Ge-based on-chip laser sources and propose that on-chip laser sources that enable rapid power gating need to be developed to enable the adoption of photonic NoC in mainstream systems. Similarly, in [8], the authors use Clos and crossbar topologies to make an argument for using on-chip laser sources in manycore NoC from an energy-efficiency and energy-proportionality perspective.

Several design flows have been developed to enable rapid exploration of the silicon-photonic link design space as well as the silicon-photonic NoC design space. In [20], the authors propose a linear programming technique to design the physical layout of the photonic devices on a separate photonic layer with the goal of minimizing the power consumed by an off-chip laser source. Similarly, in [21], the authors propose a methodology to route the photonic waveguide such that the number of waveguide crossings is minimized. For a comprehensive evaluation of the photonic network design space, the authors in [22] propose a methodology and a tool that jointly explores the link-level and system-level designs of the network topologies. In [23], the authors propose a design flow for placement and routing of photonic devices to hierarchically design large complex photonic networks. In [24], the authors have proposed a tool for placement and routing of optical NoC topologies with the goal of enabling a realistic analysis of the optical NoC design space.

In our paper we explore the design space for sharing and placement of on-chip laser sources. Our methodology simultaneously considers NoC bandwidth constraints, thermal constraints and physical layout constraints to determine on-chip laser source sharing and placement that reduce laser power consumption. This methodology can be readily applied to any of the on-chip laser source technologies under development. It is also possible to integrate our design strategy with design automation approaches that focus on optimizing other aspects of silicon-photonic NoC.

### III. EXPERIMENTAL SETUP

To explore the limits and opportunities for sharing and placement of laser sources, we consider a 3D stacked manycore

system (see Fig. 1(a)) with a logic layer containing 256 cores fabricated using standard bulk CMOS process, a photonic NoC layer next to the metal stack and a layer for on-chip laser sources. The logic layer and photonic layer are connected using vertical metal vias. The architecture of each core in the logic layer is similar to an IA-32 core used in the Intel Single-Chip Cloud Computer [25]. We scale the core power and dimensions from 45 nm to 22 nm technology [26], resulting in a total chip area of 366.1 mm<sup>2</sup> (0.93 mm<sup>2</sup> per core, including L1, and 0.50 mm<sup>2</sup> for each 256 KB private L2 cache). We choose the operating frequency as 800 MHz and the supply voltage as 0.65 V, and scale the per-core power based on the reported data of Intel 22 nm Tri-Gate technology. The average per core power is 0.46 W, and the average per L2 cache power is 0.01 W (based on ITRS-LSTP cache model in CACTI [27]). There are 16 memory controllers uniformly distributed along the two edges of the chip. The silicon-photonic NoC in our system is used for connecting the private L2 caches of the cores and the memory controllers. It should be noted that our proposed methodology for sharing and placement of laser sources is also applicable for a shared L2 architecture where the photonic NoC provides connectivity between L1 cache and shared L2 cache.

For evaluation, we use Sniper [28] simulator for system-level simulations and then use McPAT [29] to derive dynamic core power values. We select a representative set of multi-threaded benchmarks from SPLASH2 [30] (*barnes*, *ocean*, *radix*, *lu\_contiguous*, *fft* and *water\_nsquare*) and PARSEC [31] (*blackscholes*, *canneal* and *swaptions*) suite. We run each benchmark with *sim\_medium* input and 256 threads, and then calibrate the core power consumption based on the average power consumption given above. As for the bisection bandwidth of target system, we choose 512 GB/s based on *canneal*, which has the highest bandwidth requirements among all benchmarks we evaluated. Based on the calculated core power consumption when running the above listed benchmarks, we select 0.4 W and 0.7 W as the lower bound and upper bound, respectively, for core power consumption in the following case studies.

For our target system, light waves emitted by one or more single-mode indium phosphide (InP)-based laser sources located on the laser layer are routed into silicon (Si) waveguides located on the photonic device layer. Instead of single-mode lasers, integrated multi-wavelength lasers [32] or comb

lasers [33] could be utilized. However, these alternatives have a large footprint or require further technological development. Therefore only single-mode laser sources are considered in this work. We place the single-mode laser sources over the lower power-density L2 cache blocks to minimize the impact of the core power on the laser source temperature. The laser source is driven by a driver located in the logic layer. With the exception of the photodetector which is made of germanium ( $Ge$ ), all other photonic devices in the photonic device layer are made of  $Si$ . The cladding material is silicon dioxide ( $SiO_2$ ). The light waves pass next to a ring modulator that is driven through the metal vias by the link transmitter circuit located on the logic layer. The modulators convert data from the electrical domain to the photonic domain. The modulated light waves propagate in the waveguides and can pass through zero or more ring filters. At the receiver side, the light waves are filtered by ring filters and then are absorbed by  $Ge$  photodetectors. The current generated by the photodetectors passes through the metal vias and is fed to the link receiver circuit located on the logic layer.

To explain the various tradeoffs associated with choosing the laser source configuration across various logical topologies, we compare an 8-ary 3-stage Clos topology, a 16-ary 3-stage Clos topology and a  $16 \times 16$  crossbar mapped to a U-shaped physical layout of the waveguides in the photonic layer of our target system. The choice of these topologies is driven by the fact that silicon-photonic link technology is most appropriate for high-radix low-diameter topologies like Clos and crossbar. We also investigate the tradeoffs associated with choosing the laser source configuration when the  $16 \times 16$  crossbar and the 16-ary 3-stage Clos are mapped to U-shaped and W-shaped layouts shown in Fig. 1(a). For our analysis, we use the projected photonic losses listed in Table I. It should be noted that our proposed methodology for choosing a laser source is generally applicable to any physical layout and any logical topology mapped to that layout.

To evaluate the impact of temperature variations (due to variations in core power and laser power) on the laser efficiency, we use the 3D extension of HotSpot 5.02 [34] for our thermal simulations. We set the ambient temperature to  $35^\circ C$  and use the default package configurations in HotSpot. The cross-sectional view of the 3D system that we evaluated is shown in Fig. 1(b). For thermal analysis, the laser sources are modeled individually on the laser layer. However, for waveguides and rings on the photonic layer, we aggregated them into larger-sized blocks in the floorplan as using a separate model for every waveguide and ring leads to large simulation time without any significant improvement in precision. Our aggregation methods provide desirable accuracy-simulation time tradeoffs in thermal simulation [35]. We compute the joint thermal resistivity for waveguide blocks and ring blocks using  $R_{joint} = V_{total}/\Sigma(V_i/R_i)$ , where  $R_i$  and  $V_i$  refer to the thermal resistivity and volume of material  $i$  in the blocks. The dimensions of our system are shown in Fig. 1(b). All the thermal results we report in this work are from steady state analysis.

#### IV. ON-CHIP LASER SOURCE OPTIMIZATION

This section describes the laser source model used to evaluate the laser source operating regimes. A discussion of the optical power, efficiency and thermal tradeoffs is provided,

TABLE I: **Optical Loss per Component [2]**

Photonic device	Optical Loss (dB)
Coupler	1
Splitter	0.2
Non-linearity (at 30 mW)	1
Modulator Insertion	0
Waveguide (per cm)	1~5
Waveguide crossing	0.05
Filter through	1e-4
Filter drop	1.5
Photodetector	0.1

followed by a description of the methodology to determine the sharing and placement of on-chip laser sources for minimizing laser power consumption.

##### A. Laser source model

The laser wall-plug efficiency ( $\eta_{WPE}$ ) is given by the ratio of optical output power ( $P_o$ ) to electrical input power ( $P_{IN}$ ):

$$\eta_{WPE} = \frac{P_o}{P_{IN}}, \quad (1)$$

where  $P_o$  is given by

$$P_o = \eta_i \eta_d \frac{hc}{\lambda q} (I - I_{th}), \quad (2)$$

where  $\eta_i$  and  $\eta_d$  are the laser internal efficiency and differential quantum efficiency, respectively;  $h$ ,  $c$  and  $q$  are Planck's constant, the speed of light and the elementary electric charge, respectively;  $\lambda$  is the laser operating wavelength;  $I$  and  $I_{th}$  are the drive and threshold currents, respectively [36].

The electrical input power of the laser is the product of the drive current and the total voltage across the laser's terminals and can be calculated as:

$$P_{IN} = I^2 R_s + IV_d, \quad (3)$$

where  $R_s$  is the laser series resistance and  $V_d$  represents the diode voltage.

One of the shortcomings of semiconductor lasers is the strong dependence of  $P_o$  on temperature. Fortunately, simple empirical formulae match quite well with the measured characteristics of many different semiconductor diode lasers [36]. These empirical formulae are:

$$I_{th} = I_{th}^0 e^{T/T_0}, \quad (4a)$$

$$\eta_d = \eta_d^0 e^{-T/T_\eta}, \quad (4b)$$

where  $T_0$  and  $T_\eta$  are the characteristic temperatures of the threshold current and the differential quantum efficiency, respectively, while  $I_{th}^0$  and  $\eta_d^0$  are the threshold current and the differential quantum efficiency projected to a reference temperature. Additionally, the diode voltage  $V_d$  depends on temperature through the Shockley diode equation:

$$V_d = \frac{k_B T}{q} \ln \left( \frac{I}{I_s} \right), \quad (5)$$

where  $k_B$  is the Boltzmann constant and  $I_s$  is the reverse bias saturation current. By substituting Eqs. (4) and (5) into Eqs. (2) and (3), simple relationships are expressed for the dependence




Fig. 2: (a) P-I characteristics of a laser source at various temperatures (b) Wall-plug efficiency vs. Input current at various temperatures (c) Wall-plug efficiency vs. Laser Source Lengths at various temperature.




Fig. 3: Laser source temperature vs. electrical input power for a  $300 \mu m \times 50 \mu m$  laser source.

of the laser performance on operating temperature. This model is well established and the parameters used in simulations are extracted from measurement results [37], [38]. In this paper, we consider a strained  $InP$ -based multi-quantum well (MQW) laser structure.

### B. Laser source optical power, wall-plug efficiency and temperature tradeoffs

Figure 2(a) presents the optical output power of the laser source versus the input current (P-I characteristic) for various temperatures and demonstrates that the threshold current  $I_{th}$  increases with temperature while the laser optical output power decreases with increase in temperature for a given current. Figure 2(b) demonstrates that at a given temperature, the efficiency of the laser sources initially increases, reaches a peak value and then decreases as the input current is increased. The peak efficiency decreases at higher temperature as expected from the model presented in the previous section (Eqs. 4). It is therefore desirable to operate the laser source at a low temperature and ensure that the input current is such that the laser efficiency is maximized. Figure 2(c) shows the variation of laser source efficiency with laser source length (for lasers operating at the optimal current level) for various laser source temperatures. Here the laser source width is fixed at  $50 \mu m$  while the laser source length is varied from  $200 \mu m$  to  $700 \mu m$ . Figure 2(c) shows that a laser source that is  $300 \mu m$  long has the highest efficiency at any given temperature. This is because for shorter cavity length, high-order effects result in a reduction of the carrier density above threshold, which in turn decreases  $\eta_i$ . For longer cavity length,  $\eta_d$  dominates and the efficiency decreases. This behavior is typical of the strained  $InP$  laser structures that are used for our simulations. Therefore a laser source that is  $300 \mu m \times 50 \mu m$  is used for the remaining analysis.

To determine the impact of electrical input power of the laser sources on laser source temperature, we ran Hotspot simulations for the 256-core target system with each core consuming  $0.4 W$ ,  $0.5 W$ ,  $0.6 W$  and  $0.7 W$  of power. As shown in Fig. 3, as the laser input power increases the laser source temperature increases, which reduces the laser source efficiency. The increase in the core power consumption in the logical layer also increases the temperature of the laser source, which further lowers the laser source efficiency.

Based on the power-temperature-efficiency tradeoffs of the laser source shown above, Fig. 4 shows the laser source efficiency and electrical input power as a function of optical output power of the laser source for two systems – one where each core is dissipating  $0.4 W$  of power and the second where each core is dissipating  $0.7 W$  of power. The required optical output power from a laser depends on the optical losses in the photonic link being driven by that laser source. For the case where each core's power consumption is  $0.4 W$ , Fig. 4(a) shows that the optimal operation point is for a laser output power of  $23 mW$  per wavelength, where the laser source achieves a peak efficiency of 8.2%. This results in a laser electrical input power of  $268 mW$ . When each core's power consumption is  $0.7 W$ , the optimal laser output power is still approximately  $23 mW$ , but the laser efficiency decreases to 6.2% due to the higher laser temperature. This results in an electrical input power of  $355 mW$ . This analysis demonstrates that the electrical input power of the laser increases with increase in core power consumption. Hence, one needs to develop run-time techniques that can minimize core and laser power to minimize the peak temperature of the laser and in turn maximize laser efficiency and minimize the laser input power.

As shown in Fig. 4, a laser source outputs a specific optical output power at its maximum efficiency. Depending on the optical power required per  $\lambda$  and the laser technology used, it may be desirable to share laser source power across two or more waveguides. Figure 5 shows two methods for sharing single-mode laser sources across multiple waveguides. Figure 5(a) uses ring filters at the crossing of waveguides to filter and route each wavelength into the waveguide such that each waveguide propagates a number of wavelengths. Assuming each waveguide crossing incurs  $0.05 dB$  loss,  $1e-4 dB$  through loss per ring [39] and 64 waveguides share each laser source, this sharing approach introduces an overhead of  $3.2 dB$  of optical loss. Figure 5(b) shows an alternate method that first couples the light waves from a set of laser sources (each emitting one wavelength), and then splits the light into




Fig. 4: Wall-plug efficiency vs. Optical output power by the laser source for different granularity of sharing while a background logical layer operates at 0.4 W per core (a) and 0.7 W per core (b) respectively.




Fig. 5: (a) The laser source sharing through ring filters at the crossings of waveguides (b) The laser source sharing through waveguide couplers and splitters.

multiple waveguides. Assuming a loss of 0.2 dB/split, if 64 waveguides are sharing these laser sources, this sharing method causes an overhead of 1.2 dB optical loss that is lower than the overhead of the other method. Therefore, we choose the latter sharing method in this paper.

Figure 6 shows a plot of variation in laser efficiency with optical output power per wavelength for different granularity of sharing a single-mode laser source across multiple waveguides. We consider two cases with core power of 0.4 W and 0.7 W. For the 0.4 W case, if the total optical power required per wavelength is small, say 1 mW, then using a laser source per wavelength per waveguide results in a laser source efficiency of 1%. If we were to share the laser source across two waveguides, the total required optical output power of the laser source increases, which increases the efficiency to 2%. For this 1 mW optical power per wavelength per waveguide case, it makes most sense to share the laser source across 16 waveguides as it provides the maximum efficiency. Sharing of laser sources across more than 16 waveguides increases the sharing cost and hence the laser source needs to emit larger optical power, which decreases the laser source efficiency. Thus, an optimal sharing of laser sources is critical for operating laser sources at maximum efficiency and minimizing the electrical input power. A similar argument can be made for the case where the core power is 0.7 W. The only difference is that in this case the laser source efficiency is lower due to the higher core power.

### C. On-chip laser source sharing and placement strategy

To determine the sharing and placement of the laser source, we propose a cross-layer approach where we jointly consider the NoC bandwidth constraints driven by the applications running on the manycore system, thermal constraints driven by the power consumed by the cores and the laser sources, physical layout constraints driven by the losses in the photonic devices and laser source designs that are compatible with our

proposed 3D system. Figure 7 shows a flowchart describing the strategy for determining the sharing and placement of laser sources in the 3D system.

The number of cores in the target system and the type of applications that are expected to run on the target system can be used to determine the amount of traffic that will be injected into the NoC. This in turn can be used to determine the optimal NoC topology at the logic level and the bandwidth of each channel in the NoC. Depending on the bandwidth per channel, operating frequency of cores and bandwidth of each silicon-photonic link, we can determine the number of silicon-photonic links, i.e. number of wavelengths required by the target system. The chosen logical topology for an NoC can be mapped to several different layouts [40]. Depending on the loss components (including waveguide loss, through loss, crossing loss, etc.) in a silicon-photonic link, we identify various potential physical layouts of the NoC with three candidates for the placement and sharing of the laser source – 1) all laser sources are placed locally next to the router with a laser source emitting one wavelength for one waveguide (no sharing); 2) all laser sources are placed locally next to the router with each laser source emitting one wavelength that is shared across two or more waveguides; and 3) all laser sources are placed along the edge with each laser source emitting one wavelength shared by two or more waveguides. For these candidate layouts and the target bandwidths per channel, depending on the thermal properties, i.e. expected peak temperature, of the laser source at runtime and the laser source efficiency characteristics, we can decide the placement of the laser source and its sharing granularity such that the total laser power consumption is minimized.

## V. CASE STUDIES

In this section, we present two case studies to show how the sharing and placement of laser sources change with logical topologies and physical layouts. For this analysis we assume the target system bisection bandwidth is always 512 GB/s, all cores are consuming 0.7 W of power and we use the results in Fig. 4(b) and 6(b) to determine the laser source efficiency.

### A. Laser source placement and sharing across various logical topologies

Figure 8(a) shows the total electrical input power of the laser sources for an 8-ary 3-stage Clos topology for different placements of laser sources. This 8-ary 3-stage Clos network uses 24 routers (8 routers in each stage), and each router in

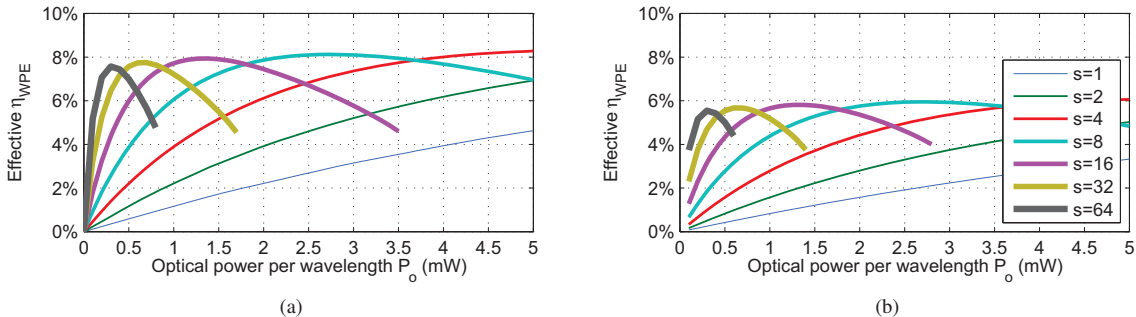


Fig. 6: The effective WPE while increasing the sharing of laser sources as the logical layer operates at 0.4 W (a) and 0.7 W (b) per core, respectively.  $s = x$  indicates that  $x$  waveguides are sharing the output of the laser source. Splitter loss is accounted.

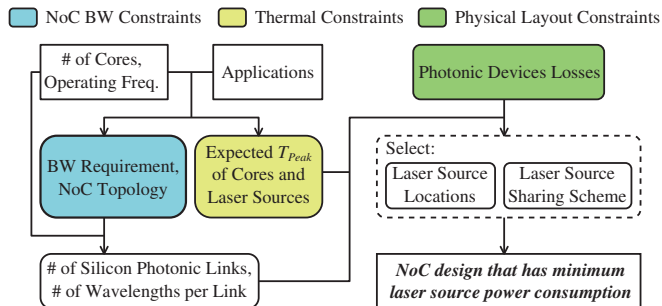


Fig. 7: Flowchart for deciding the sharing and placement configurations of on-chip laser sources.

the 1<sup>st</sup> and 3<sup>rd</sup> stage is connected to 32 cores (4 cores per router input). There are 64 photonic channels connecting the 1<sup>st</sup> and 2<sup>nd</sup> stage of routers and another 64 connecting the 2<sup>nd</sup> and 3<sup>rd</sup> stage of routers. The detailed specifications of the 8-ary 3-stage Clos are given in Table II. The photonic channels are mapped to a U-shaped layout shown in Fig. 1(a). In the physical layout of this 8-ary 3-stage Clos topology, one router each from the 1<sup>st</sup> stage, 2<sup>nd</sup> stage and 3<sup>rd</sup> stage are placed next to each other. For a typical waveguide loss of 2 dB/cm, 0.15 mW of optical output power per wavelength is required. If local laser sources are not shared, the efficiency of each local laser source for 0.15 mW optical output power is 0.12%. This results in a total electrical input power for laser sources of 243 W (119 mW per laser source).

Given that the routers in the 1<sup>st</sup>, 2<sup>nd</sup> and 3<sup>rd</sup> stage of the Clos network are placed next to each other, there is an opportunity for sharing the local laser sources among the 16 photonic channels (8 from the 1<sup>st</sup> stage router and 8 from the 2<sup>nd</sup> stage router), whereby the optical output power of a laser source is split and routed into the waveguides associated with these photonic channels. This sharing of the local laser sources increases the total optical output power of each laser source, which improves its efficiency. For this particular example, each one of the 16 photonic channels is mapped to a waveguide with 16  $\lambda$  per photonic channel, i.e. 16  $\lambda$  per waveguide. If each laser source is shared across these 16 waveguides, the total optical output power is 2.4 mW for each  $\lambda$ , which corresponds to a laser source efficiency of 1.3% and a total electrical input power of 23.63 W (185 mW per laser source).

For the laser source example considered in this paper, the maximum efficiency is achieved at an optical output power of 23 mW. If we want to use a laser source that outputs 23 mW of power, then we propose to place the laser source

along the edge so that there are more opportunities for sharing. The 128 photonic channels in the 8-ary 3-stage Clos topology correspond to 128 waveguides. For the case where the laser sources are placed along the edge, the optical output power required for each waveguide is 0.18 mW for each  $\lambda$ . This value is higher than 0.15 mW due to longer waveguide lengths. We can share 16 laser sources (1 for each  $\lambda$ ) across these 128 waveguides so that they can operate at 6.38% efficiency. This results in a total electrical input power of the lasers of 5.74 W (359 mW per laser source).

As the waveguide loss per cm increases, the total optical output power required increases. In the case of non-shared local laser sources, this increment in optical output power improves the efficiency. As a result, the total electrical input power does not increase significantly. In the case of shared local laser sources, the increase in the optical power requirement per waveguide pushes the efficiency of laser sources towards the peak value. Hence, similar to non-shared local laser sources, there is only a marginal increment in the electrical input power of the laser sources. The layout with laser sources located along the edge has longer waveguides, and so the optical loss increases significantly when waveguide loss per cm increases. Hence, the required optical power per waveguide increases, which in turn lowers the laser source efficiency. Overall, if the waveguide loss is low (< 3 dB/cm), using shared laser sources located along the edge is the best option. On the other hand, if the waveguide loss is high (> 3 dB/cm), placing shared laser sources locally is beneficial.

For the same 256-core target system, we could use a 16-ary 3-stage Clos network for less contention among cores at the input of each router. This 16-ary 3-stage Clos topology has 48 routers (16 routers in each stage) with each router in the 1<sup>st</sup> and 3<sup>rd</sup> stage connected to 16 cores (1 core per router input). This network topology requires a total of 512 channels. To match the bisection bandwidth of this topology with the 8-ary 3-stage Clos, each channel needs 4  $\lambda$ , and the system has a total of 128 waveguides with 4 channels (4  $\lambda$  for each channel) sharing one waveguide (16  $\lambda$  in each waveguide). In general, the trends for the electrical input power of the laser for the 16-ary 3-stage Clos are similar to the trends for the 8-ary 3-stage Clos. One exception is that the electrical input power for the case using shared local laser sources is higher for the 16-ary 3-stage topology due to the decrease in the degree of sharing of laser sources.

We could also use a 16  $\times$  16 Single-Write-Multiple-Read (SWMR) crossbar topology having a concentration of 16, i.e.

TABLE II: Architectural-level parameters for 5 NoCs under consideration. U-shaped and W-shaped layouts are shown in Fig. 1.

Logical topology	Physical layout	Dimension	Concentration	$\lambda$ /Channel	Number of channels
Clos	U-shaped	8-ary 3-stage	4	16	128
Clos	U-shaped	16-ary 3-stage	1	4	512
Crossbar	U-shaped	$16 \times 16$	4	64	32
Clos	W-shaped	16-ary 3-stage	1	4	512
Crossbar	W-shaped	$16 \times 16$	4	64	32

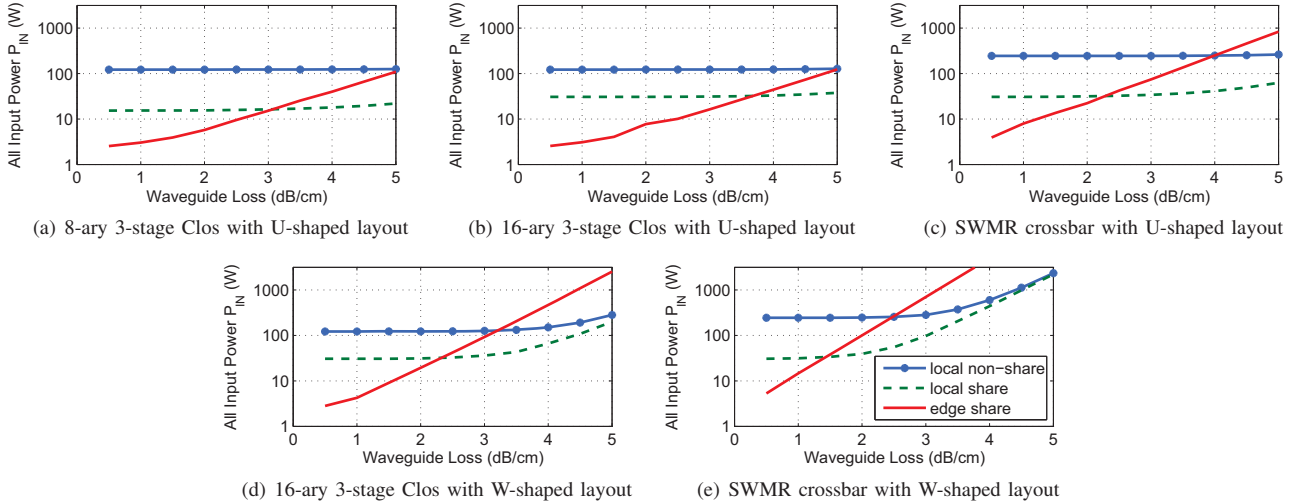


Fig. 8: Total laser power vs. waveguide loss for various sharing scenarios and placements of on-chip laser sources. (a), (b) and (c) compare various topologies with U-shaped layout. (b) and (d) compare various layouts for 16-ary 3-stage Clos topology. (c) and (e) compare various layouts for SWMR crossbar topology. We assume each core in the logical layer operates at  $0.7 W$ .

each input of the crossbar can be accessed by 16 cores (similar to 16-ary 3-stage Clos) for the NoC and map it to the U-shaped physical layout. In this case, to match the bisection bandwidth of the crossbar with the Clos networks, we need  $64 \lambda$  for each channel (4 waveguides per channel). Figure 8(c) shows the total electrical input power for the lasers. Compared to the Clos network, the SWMR channels are shared by more receivers, therefore the large number of rings along longer waveguides results in higher laser power consumption than the Clos networks. For smaller waveguide losses ( $< 2 \text{ dB/cm}$ ), shared laser sources located along the edge are preferable. On the other hand, for larger waveguide losses ( $> 2 \text{ dB/cm}$ ), using shared local laser sources are preferable.

Overall, the best sharing and placement of on-chip laser sources depend on the network topology. For example, at  $3.5 \text{ dB/cm}$  waveguide loss, for 8-ary 3-stage Clos and SWMR crossbar mapped to U-shape physical layout, using shared local laser sources minimizes the electrical input power, while for 16-ary 3-stage Clos using shared laser sources located along the edge is the better choice. On the other hand for a  $2 \text{ dB/cm}$  waveguide loss, shared local laser sources are preferable for all three logical topologies.

### B. Laser source placement and sharing across various physical layouts

Depending on alternate power, performance and area design constraints, the placement and routing tools may generate physical layouts that are different from the U-shaped layout that we considered in the earlier subsection. The choice of

laser source sharing and placement changes with a change in the physical layout. For example, mapping the 16-ary 3-stage Clos topology and the  $16 \times 16$  SWMR crossbar to a W-shaped layout (see Fig. 1(a)) increases optical waveguide losses, and hence the electrical input power required for a laser source. Figure 8(d) and 8(e) shows the total electrical input power for the lasers varying with waveguide loss per  $cm$  for the 16-ary 3-stage Clos and  $16 \times 16$  SWMR crossbar with W-shaped layout, respectively. Similar to the U-shaped layout, for low waveguide loss shared laser sources placed along the edge are preferable, while for high waveguide loss shared local laser sources are preferable. The crossover point (i.e. where the choice of laser source placement and sharing changes from shared laser sources placed along the edge to shared laser source placed locally) is different for the two layouts. For the 16-ary 3-stage Clos topology, the crossover points are  $3.6 \text{ dB/cm}$  and  $2.2 \text{ dB/cm}$  for the U-shaped and W-shaped layouts, respectively. On the other hand, for the  $16 \times 16$  SWMR crossbar topology, the crossover points are  $2.2 \text{ dB/cm}$  and  $1.5 \text{ dB/cm}$  for the U-shaped and W-shaped layouts, respectively.

## VI. CONCLUSION

In this paper, we explored limits and opportunities for sharing and placement of on-chip laser sources that drive the silicon-photonics NoC with the goal of minimizing the total laser power consumption. We first explored the power, efficiency and temperature tradeoffs associated with on-chip laser sources. Then we proposed a cross-layer methodology

that jointly considers NoC bandwidth constraints, thermal constraints and physical layout constraints to determine the sharing and placement of laser sources. We explored the application of our methodology to a 256-core system by considering three different logical topologies, two different physical layouts and three different sharing/placement strategies for its NoC. Our analysis shows that the choice of laser source placement and sharing changes with the choice of logical topology, physical layout and waveguide loss. For a 8-ary 3-stage Clos mapped to a U-shaped layout, with waveguide loss lower than  $3 \text{ dB/cm}$ , shared laser sources placed along the edge consume the least laser power, while at waveguide loss higher than  $3 \text{ dB/cm}$ , local shared laser sources provide the least laser power. For a  $16 \times 16$  SWMR crossbar topology with matching bisection bandwidth and mapped to the U-shaped layout, shared laser sources placed along the edge are preferable for waveguide losses less than  $2.2 \text{ dB/cm}$ , while local shared laser sources are preferable for waveguide losses greater than  $2.2 \text{ dB/cm}$ . For the same crossbar topology mapped to a W-shaped layout, the shared laser sources placed along the edge are preferable for waveguide losses less than  $1.5 \text{ dB/cm}$ , while local shared laser sources are preferable for waveguide losses greater than  $1.5 \text{ dB/cm}$ .

#### ACKNOWLEDGEMENT

This work has been partially supported by the NSF grants CCF-1149549 and CNS-1149703.

#### REFERENCES

- [1] N. Kirman and J. F. Martinez, "A power-efficient all-optical on-chip interconnect using wavelength-based oblivious routing," in *ASPLOS*, 2010.
- [2] A. Joshi *et al.*, "Silicon-photonics clos networks for global on-chip communication," in *NOCS*, 2009.
- [3] D. Vantrease *et al.*, "Corona: System implications of emerging nanophotonic technology," in *Proc. ISCA*, ser. ISCA '08, 2008, pp. 153–164.
- [4] Y. Pan *et al.*, "Firefly: illuminating future network-on-chip with nanophotonics," in *ISCA*, 2009.
- [5] M. J. Cianchetti, J. C. Kerekes, and D. H. Albonesi, "Phastlane: a rapid transit optical routing network," in *ISCA*, 2009.
- [6] A. Shacham, K. Bergman, and L. P. Carloni, "On the design of a photonic network-on-chip," in *NOCS*, 2007, pp. 53–64.
- [7] G. Kurian *et al.*, "Cross-layer energy and performance evaluation of a nanophotonic manycore processor system using real application workloads," in *Proc. IPDPS*, May 2012, pp. 1117–1130.
- [8] M. Heck and J. Bowers, "Energy efficient and energy proportional optical interconnects for multi-core processors: Driving the need for on-chip sources," *IEEE JSTQE*, vol. 20, no. 4, pp. 1–12, July 2014.
- [9] C. Chen and A. Joshi, "Runtime management of laser power in silicon-photonics multibus noc architecture," *IEEE JSTQE*, vol. 19, no. 2, pp. 3 700 713–3 700 713, 2013.
- [10] C. Li *et al.*, "Luminoc: A power-efficient, high-performance, photonic network-on-chip for future parallel architectures," in *NOCS*, 2013.
- [11] L. Zhou and A. K. Kodi, "PROBE: Prediction-based optical bandwidth scaling for energy-efficient NoCs," in *NOCS*, 2013.
- [12] T. Zhang *et al.*, "Thermal management of manycore systems with silicon-photonics networks," in *Proc. DATE*, 2014.
- [13] C. Nitta, M. Farrens, and V. Akella, "Addressing system-level trimming issues in on-chip nanophotonic networks," in *Proc. HPCA*, 2011, pp. 122–131.
- [14] Y. Zheng *et al.*, "Power-efficient calibration and reconfiguration for on-chip optical communication," in *DATE*, 2012, pp. 1501–1506.
- [15] Z. Li *et al.*, "Reliability modeling and management of nanophotonic on-chip networks," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 20, no. 1, pp. 98–111, 2012.
- [16] A. W. Fang *et al.*, "Electrically pumped hybrid alginas-silicon evanescent laser," *Opt. Express*, vol. 14, no. 20, pp. 9203–9210, Oct 2006.
- [17] R. E. Camacho-Aguilera *et al.*, "An electrically pumped germanium laser," *Opt. Express*, vol. 20, no. 10, pp. 11 316–11 320, May 2012.
- [18] L. Liu *et al.*, "A thermally tunable iii-v compound semiconductor microdisk laser integrated on silicon-on-insulator circuits," *IEEE Photonics Technology Letters*, vol. 22, no. 17, pp. 1270–1272, Sept 2010.
- [19] T. Wang *et al.*, "1.3- $\mu\text{m}$  inas/gaas quantum-dot lasers monolithically grown on si substrates," *Opt. Express*, vol. 19, no. 12, pp. 11 381–11 386, Jun 2011.
- [20] D. Ding *et al.*, "O-router: An optical routing framework for low power on-chip silicon nano-photonics integration," in *Proc. DAC*, July 2009, pp. 264–269.
- [21] C. Condrat, P. Kalla, and S. Blair, "Crossing-aware channel routing for photonic waveguides," in *Proc. MWSAS*, Aug 2013, pp. 649–652.
- [22] J. Chan *et al.*, "Physical-layer modeling and system-level design of chip-scale photonic interconnection networks," *Trans. Comp.-Aided Des. Integ. Cir. Sys.*, vol. 30, no. 10, pp. 1507–1520, Oct. 2011.
- [23] G. Hendry *et al.*, "Vandal: A tool for the design specification of nanophotonic networks," in *Proc. DATE*, 2011, pp. 1–6.
- [24] A. Boos *et al.*, "Proton: An automatic place-and-route tool for optical networks-on-chip," in *Proc. ICCAD*, Nov 2013, pp. 138–145.
- [25] J. Howard *et al.*, "A 48-core IA-32 processor in 45 nm CMOS using on-die message-passing and DVFS for performance and power scaling," *IEEE Journal of Solid-State Circuits*, vol. 46, no. 1, pp. 173–183, 2011.
- [26] J. Meng *et al.*, "Run-time energy management of manycore systems through reconfigurable interconnects," in *Proc. GLSVLSI '11*, 2011, pp. 43–48.
- [27] S. Thoziyoor *et al.*, "CACTI 5.1," HP Laboratories, Palo Alto, Tech. Rep., 2008.
- [28] T. E. Carlson, W. Heirman, and L. Eeckhout, "Sniper: Exploring the level of abstraction for scalable and accurate parallel multi-core simulations," in *Proc. SC*, 2011, pp. 1–12.
- [29] S. Li *et al.*, "McPAT: An integrated power, area, and timing modeling framework for multicore and manycore architectures," in *Proc. MICRO-42*, 2009, pp. 469–480.
- [30] S. C. Woo *et al.*, "The SPLASH-2 programs: characterization and methodological considerations," in *Proc. ISCA*, 1995, pp. 24–36.
- [31] C. Bienia *et al.*, "The PARSEC benchmark suite: Characterization and Architectural Implications," in *Proc. PACT*, 2008, pp. 72–81.
- [32] K. Lawniczuk *et al.*, "8-channel awg-based multiwavelength laser fabricated in a multi-project wafer run," in *Proc. International Conference on Indium Phosphide and Related Materials*, May 2011, pp. 1–4.
- [33] D. Livshits *et al.*, "Cost-effective wdm optical interconnects enabled by quantum dot comb lasers," in *Proc. SPIE*, 2010, pp. 1–9.
- [34] J. Meng, K. Kawakami, and A. K. Coskun, "Optimizing energy efficiency of 3-D multicore systems with stacked dram under power and thermal constraints," in *DAC*, 2012.
- [35] A. K. Coskun *et al.*, "Modeling and dynamic management of 3D multicore systems with liquid cooling," in *VLSI-SoC*, 2009.
- [36] L. Coldren *et al.*, *Diode Laser and Photonic Integrated Circuits*, 2nd ed. Wiley Series in Microwave and Optical Engineering, 2012.
- [37] S. Hu *et al.*, "High-efficiency and low-threshold ingaas/algaas quantum-well lasers," *Journal of Applied Physics*, vol. 76, no. 6, pp. 3932–3934, Sep 1994.
- [38] —, "Lateral carrier diffusion and surface recombination in ingaas/algaas quantum-well ridge waveguide lasers," *Journal of Applied Physics*, vol. 76, no. 8, pp. 4479–4487, Oct 1994.
- [39] C. Batten *et al.*, "Building Manycore Processor-to-DRAM Networks with Monolithic CMOS Silicon Photonics," *IEEE Micro*, vol. 29, no. 4, pp. 8–21, 2009.
- [40] —, "Designing chip-level nanophotonic interconnection networks," in *Integrated Optical Interconnect Architectures for Embedded Systems*. Springer New York, 2013, pp. 81–135.