mNoC: Large Nanophotonic Network-on-Chip Crossbars with Molecular Scale Devices

Technical Report CS-2013-02

Jun Pang

Department of Computer Science
Duke University
pangjun@cs.duke.edu

Chris Dwyer

Department of Electrical and Computer Engineering
Duke University
dwyer@ece.duke.edu

Alvin R. Lebeck

Department of Computer Science
Duke University
alvy@cs.duke.edu

Abstract

Moore's law and the continuity of device scaling have led to an increasing number of cores/nodes on a chip, creating a need for new mechanisms to achieve high-performance and power-efficient Network-on-Chip (NoC). Nanophotonics based NoCs provide for higher bandwidth and more power efficient designs than electronic networks. Present approaches often use an external laser source, ring resonators for signal modulation and filtering, and waveguides for transmission. However, they still suffer from important limitations: large static power consumption, and limited network scalability.

In this paper, we explore the use of emerging molecular scale devices to construct nanophotonic networks — called Molecular-scale Network-on-Chip (mNoC). We leverage quantum dot LEDs, which provide electrical to optical signal modulation, and chromophores, which provide optical signal filtering for receivers. These devices replace the ring resonators and the external laser source used in contemporary nanophotonic NoCs. They reduce energy consumption or enable scaling to larger crossbars for a reduced energy budget. We present a Single Writer Multiple Reader (SWMR) bus based crossbar mNoC. Our evaluation shows that an mNoC can achieve more than 77% *reduction in energy for a 64x64 crossbar compared to similar ring resonator based designs. Additionally, an mNoC can scale to a 256x256 crossbar with an average* 10% *performance improvement and* 49% *energy reduction. We also discuss implications of the new mNoC crossbar on overall system design. For example, a large single crossbar allows for the possibility of high radix routers and efficient broadcast based directory protocols.*

1. Introduction

Today's microprocessor chips incorporate an increasing number of cores/nodes. To support communication among the many nodes, an on-chip network (often called Network-on-Chip or NoC) must meet various design targets, such as latency, bandwidth, area and power. Achieving these goals is increasingly difficult using standard CMOS technology due to power and latency limitations of metal wires, particularly for distances approaching the chip dimensions.

To overcome the challenges of CMOS wires, several studies explore (nano)photonic NoC designs. The main components of current nanophotonic NoCs are: 1) an external laser source, 2) a waveguide, 3) ring resonators for both modulation and detection. Nanophotonic NoCs exhibit superior power delay product and bandwidth compared to CMOS. Unfortunately, there still exist significant limitations to nanophotonic NoCs. Two important limitations are 1) high static power consumption due to the external laser source inefficiency and thermal ring tuning and 2)limited network scalability.

The ideal interconnect is one large crossbar that enables communication between all pairs of nodes in the absence of output conflicts. However, constructing crossbars larger than 64x64 is difficult/impractical due to the above limitations of CMOS wires or ring resonators. Connecting more than 64 nodes is possible using a multi hop network (either direct or

indirect) constructed with each intermediate router providing an NxN crossbar connectivity (N <= 64), but this introduces latency for the additional network hops.

In this paper, we explore the use of emerging molecular scale devices to construct nanophotonic networks— called Molecular-scale Network-on-Chip (mNoC). The specific molecular-scale devices we utilize are quantum dot LEDs and chromophores. The quantum dot LEDs provide electrical to optical signal modulation and the chromophores provide optical signal filtering for receivers. The chromophores replace the ring resonators and the quantum dot LED replaces external laser source used in current nanophotonic NoCs and reduce energy consumption or enable scaling to larger crossbars for a smaller energy budget. Further, these new components are easily integrated into a silicon foundry process and have been individually demonstrated, however, their use for NoCs has never been explored.

We present a Single Writer Multiple Reader (SWMR) bus-based crossbar mNoC. Our evaluation shows that an mNoC can achieve more than 77% reduction in energy for a 64x64 crossbar compared to ring resonator based designs. Additionally, for 49% less energy budget than a ring resonator 64x64 crossbar, an mNoC can scale to a 256x256 crossbar.

The properties of mNoCs can have significant implications on router micro architecture and system architecture design. With large nanophotonic crossbars it is possible to construct higher radix routers, or for systems with <=256 nodes to utilize a single crossbar. Broadcast or multicast can be efficiently supported with an mNoC without any significant change in energy consumption. This opens the possibility to support snoopy cache coherence on-chip, or to efficiently support broadcast based directory coherence protocols.

The remainder of this paper is organized as follows. Section 2 reviews ring resonator NoCs and motivates our work. The enabling technologies for mNoC are presented in Section 3 and Section 4 presents a SWMR mNoC crossbar design. We evaluate mNoC and compare energy and scalability to ring resonator NoCs in Section 5 and provide simulation results that explore the benefit of mNoCs for parallel applications by simulation with Graphite [28]. Section 6 discusses the architectural implications of an mNoC on router microachitecture and system architecture. Related work is discussed in Section 7 and we conclude in Section 8.

2. Background and Motivation

Ring resonator based NoCs are the predominate approach used in recent studies of nanophotonic on-chip networks [30, 14, 31, 38]. This section first reviews the basic operations of ring resonators and then discusses remaining limitations which motivate the work.

2.1. Ring Resonator-based NoC Overview

Nanophotonic technology is a potential solution to overcome RC delays on electrical buses. Optical signals transmit at the speed of light, which means low latency and waveguides provide near distance-independent power consumption. High bandwidth is achieved using Dense Wavelength Division Multiplexing (DWDM) to transmit up to 64 wavelengths in a single waveguide.

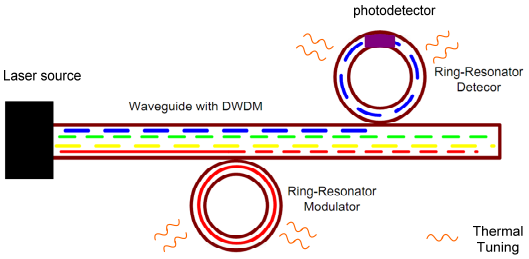

Figure 1: Main Optical Components of Ring Resonator Nanophotonic Network-on-chip

Figure 1 shows the main optical components used in current nanophotonic NoC, including a laser source, ring resonator modulators and detectors, and a waveguide. Multiple wavelengths from an external laser source are coupled to the waveguide. The ring resonator modulators are responsible for electrical to optical (E/O) conversion by modulating a specific wavelength based on electrical input. The ring resonator detector is responsible for optical to electrical (O/E) conversion. It filters out the desired wavelength and converts it to an electrical signal using a photodetector, e.g., by using a Germanium doped section on the ring. Each ring resonator must be thermally tuned (to change its refractive index) to achieve proper on/off resonance at its specified wavelength. Thermal tuning can consume 20-100μW [30, 9] per ring.

Compared with all electrical NoC designs, most current nanophotonic NoCs exhibit better performance with reduced power consumption [38, 31, 18]. However, there are at least two remaining limitations of existing nanophotonic NoC designs: large static power consumption, and limited network scalability.

2.2. Ring Resonator Limitations

First, the current nanophotonic designs are based on ring resonators for modulation and demodulation and a major limitation is the non-negligible thermal tuning power. For a single ring it is about 20μW for a 20K temperature range [14], but this incurs a large amount of power consumption in total. For example, Pan et. al [30] show that 38% of total energy for a conventional 32x32 nanophotonic crossbar is consumed by ring thermal tuning. In a radix-64 conventional nanophotonic crossbar, such as Corona [38], the thermal tuning power is around 20W for 1024K rings. When the radix increases to 128, around 84W is needed; when the radix increases to 256, over 300W is required. The prohibitively high thermal tuning power makes the crossbar impractical to scale to more than 64.

Second, most current nanophotonic designs use an activity-independent off-chip laser source, which contributes another

big portion to the total power. According to analysis from Pan et. al [30], 36% of the total energy is consumed in the electrical laser source for a radix-32 conventional nanophotonic crossbar. This is not only because the inefficiency of the off-chip laser source (30% efficiency), but also because the off-chip laser source needs to constantly couple laser power into waveguides no matter what the traffic activity is. However, the waveguide link utilization is usually very low due to the abundant on-chip resources as analyzed later in Section 5. Furthermore, since the laser source does not sit together with the sender, special designs for data channels are required [30]: two-round data channels allow the sender to modulate in the first round and the recever to demodulate in the second round; single-round data channel uses two sets of wavelengths in opposite directions. No matter which design is used, unnecessary laser power is introduced compared with an on-chip light source.

Third, the devices' nonlinearity constrains current ring resonator nanophotonic NoCs from further scaling, regardless of the previously mentioned problems. The number of devices that can be connected to a single waveguide is limited by the optical power received at the receiver photodetector after incurring maximum loss along the path from the laser source. Signal loss occurs due to waveguide loss(about 1dB/cm), insertion loss for each ring resonator on the waveguide, and other aspects such as branching/merging of waveguides. However, the maximum input optical power from the laser source cannot exceed a threshold value, due to nonlinear response of waveguides [25, 9, 6] and ring resonators [11, 34, 45], without causing incorrect operation of the respective component. Waveguide loss, nonlinearity and photodetector sensitivity appear to be sufficient for large numbers of devices; however, ring resonator nonlinearity places a severe bound on the number of devices. Recent results show that input power cannot exceed $0.6mW$ for a modulator on resonance without entering the nonlinear regime [6], while the maximum injected waveguide power without causing nonlinearity is $115mW$ [6].

Area overhead is another potential limitation on scalability. For example, a typical ring resonator used in a NoC has a radius of $5\mu m$, and a 64x64 bus-based crossbar structure requires $1024K$ ring resonators, which is at least 20% of the total area of a $400mm^2$ chip [14]. For crossbars this number grows quadratically in the number of nodes.

For the above reasons, it is difficult or infeasible to scale ring resonator NoC crossbars, and most designs have no more than 64 nodes. Scaling to a larger number of nodes requires either a multi hop network or an alternative approach. The remainder of this paper explores one potential alternative.

3. Molecular Nanophotonic Technology

This section introduces the building blocks for mNoCs and discusses their properties with respect to energy and performance. The specific molecular scale devices we utilize are quantum dot LEDs and chromophores. The quantum dot LEDs provide electrical to optical signal modulation and the chromophores

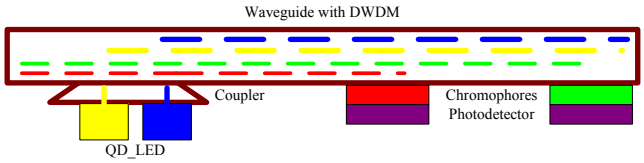

Figure 2: Main optical components of Chromophore Network-on-chip

provide optical signal filtering for receivers. These devices replace the ring resonators and the external laser source used in current nanophotonic NoCs and reduce energy consumption or enable scaling to larger crossbars for a reduced energy budget.

Figure 2 shows the main optical components of our mNoC. Quantum Dot LEDs(QD_LED) provide both the light source and modulation in a single device. All QD_LEDs from the same node inject light into the waveguide through a coupler; chromophores deposited on the waveguide work as filters and couple to a photodetector for O/E conversion. In this paper, we use commercially available chromophores that operate in the visible light range $(390-750nm)$ to filter out signals from waveguides, thus waveguides and QD_LED must also work in this range. Furthermore, to avoid crosstalk between chromophores, we limit the number of available wavelengths to only 10 instead of 64 commonly used in current nanophotonic NoCs. We also note that many quantum dots and many chromophores are used to replace a single ring resonator.

3.1. Transmitter: Quantum-Dot-LED

QD_LED light emission is controlled by the bias voltage applied across the QD_LED. By controlling the value of the applied bias voltage, the QD_LED is both a light source and modulator. QD_LED has several advantages such as small size, narrow emission bandwidth, and good stability (the output signal is resistant to degradation caused by effects such as photobleaching) [42]. Important properties to consider when using QD_LED in mNoCs include: external quantum efficiency, modulation rate, size, and energy injection rate.

External Quantum Efficiency (EQE) is defined as the number of emitted photons per injected electron and is different for QD_LEDs with different emission wavelengths. EQE determines the energy consumption of the QD_LED as a light source. EQE is 7.5% for blue light and about 20% for red light as reported by QD Vision [2]. Assuming visible light in the range from $390nm$ to $750nm$, the corresponding energy efficiency is in the range of 8.9% to 10.5%.

The modulation rate determines the number of bits that can be modulated per second. For QD_LED it is the reciprocal of the decay time of quantum dot photoluminescence. At room temperature, the decay time of QD in QD_LED is affected by the shell thickness [16] and varies from $15ps$ to $900ps$ for different shells [19, 4]. The corresponding modulation rate is

in the range $1.1GHz - 67GHz$.

By controlling the effective area of QD_LED, we can get different amounts of output power as needed. The area of QD_LED can be very small if the desired output power is small, since the size of quantum dots is from only 3 to $12nm$ in diameter, and the thickness of all the layers of QD_LED together is around $100nm$ [3]. For example, if we want to drive 255 photoreceivers with sensitivity of $0.1\mu W$ along a $18cm$ long waveguide($1dB/cm$), the output power from the QD_LED should be $345\mu W$. With 0.8 internal quantum efficiency [3], that equals to 0.2 million photons with wavelength $500nm$ and decay time $200ps$. By using contact printing fabrication [17], 10 Quantum dots can be fabricated within $100nm$ in one dimension which means the area of 100 quantum dots is $0.01\mu m^2$. For 0.2 million QDs, that is about $21.5\mu m^2$. However, if we have a $10dB$ insertion loss path which is roughly what the Corona 64×64 crossbar has [6] and drive 63 photoreceivers, then the area becomes $1.9\mu m^2$. Compared with $100\mu m^2$ area of a ring resonator, our area numbers are much smaller.

The energy injection rate corresponds to the optical energy injected into the waveguide during operation. When an off-chip laser source is used, it continuously injects light into the waveguide for ring modulators to modulate. However, if we use an on-chip QD_LED, it only injects light into the waveguide when the optical link is utilized and the transmitted signal is '1' but not '0'. Therefore, it greatly reduces the constantly coupled power from light source. If the ratio of '1' to '0' is 1, then 50% of the input power can be eliminated.

3.2. Receiver: chromophores

An mNoC utilizes chromophores to filter out optical signals and couple to photodetectors to do O/E conversion. Chromophores are small molecules that absorb light at one wavelength and emit at a different lower energy wavelength [21]. For an mNoC, signal transmission from the waveguide to the chromophores and from the chromophores to the photodetector are both through near field evanescence coupling [40, 32, 27].The chromophore couplings take near instant time, and thus the delay can be ignored.

Chromophores have different properties from ring resonators. First, they have very small size ($\sim 1nm$) as opposed to $5\mu m$ radius for a ring resonator, so there is an area range for us to choose depending on the intensity of the incident light; second, chromophores do not have nonlinear effects. Chromophore receivers have a much larger input energy range, from as low as one photon to a maximum defined by the design. We can always add more chromophores to absorb more light. Finally, chromophores do not require thermal tuning and their energy loss is proportional to the input energy.

Optical signal transmits from chromophores to waveguides and then to photodetectors through near field evanescence coupling. To prevent light from directly coupling between the waveguide and the photodetector, the photodetector must

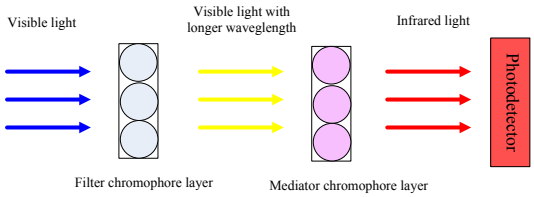

Figure 3: Receiver design

operate at a longer wavelength outside the visible range, for example in the infrared range. Germanium doped photodetectors [44, 32, 37] have good sensitivity with infrared and telecommunication wavelengths and thus are good candidates. The chromophore layers absorb visible light from the waveguide and emit light at infrared wavelength [35], which is detected by the photodetector. However, visible light from the waveguide will not be detected by the photodetector. As shown in Figure 3, we add another layer of mediator chromophores to down convert visible light energy to infrared energy through resonance energy transfer (RET). Since RET happens much faster ($9ps$) than the clock period ($200ps$), the delay can be ignored; however, the energy loss through Stoke shift [21](visible wavelengths become infrared wavelengths in our case) is considered.

3.3. Waveguide

To support mNoC, the waveguide should also operate in the visible light range. The current silicon waveguide in nanophotonic NoC [14, 18] has a pitch of $4\mu m$ with transmission loss of 1-$1.5dB/cm$. If we assume the same materials and transmission loss, in order to transmit visible light, we need a waveguide with a pitch of $2\mu m$.

The size of the visible light waveguide is about half the size of that used in current nanophotonic NoCs; however we assume only 10 instead of 64 wavelengths for one waveguide are utilized to reduce crosstalk between chromophores at the receiving side. Therefore, to support a 64-bit datapath width, at least 6 waveguides are required.

3.4. Fabrication and Experimental Demonstration

All the optical devices used in mNoC are compatible with current silicon fabrication technology [27, 36, 37, 4, 42, 19, 17]. To fabricate the receiver, the photodetector is the bottom layer with chromophores on top followed by the waveguide.

Most of the above mNoC components have been demonstrated, including QD_LEDs and waveguides in the visible light range. The only component that has not been demonstrated is the evanescent coupling of chromophores to waveguides. Previous research has demonstrated evanescently coupled photodetectors [27]. Therefore, we developed a prototype fiber-based system to demonstrate evanescent wave coupling of chromophores to a waveguide. The fiber core serves as a proxy for an on-chip visible light range waveguide. To prepare the fiber we first strip the cladding and buffer coating of a commercially available $400\mu m$ core quartz fiber, and then use dip-

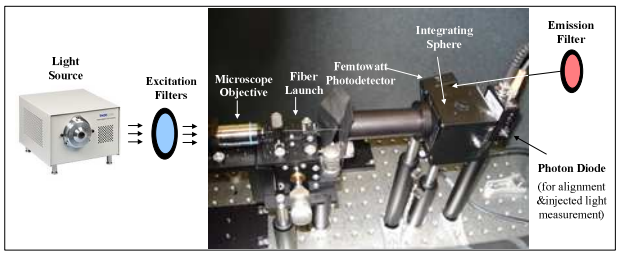

Figure 4: Experimental Setup

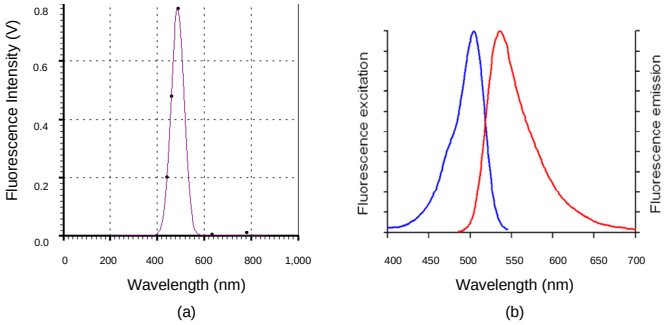

Figure 5: Excitation spectrum of Oregon Green^{TM}488 Chromophore from (a)Gaussian fit to the measured fluorescence Intensity, and (b) vendor

or drop- coating to apply one or more layers of chromophores to the exposed core.

The experimental setup to inject input light and observe output is shown in Figure 4. An incandescent broadband, i.e., white, light source passes light through Bragg interference filters to select the input wavelengths under test. The filtered light, a combination of any relevant input wavelengths, is injected into the fiber through a microscope objective aligned to focus the light onto the backend of the fiber (i.e., the side without chromophores). The active end of the fiber (with chromophores) is inserted into an integrating sphere that captures any emitted photons and focuses them on a femtowatt-sensitive detector and emission filter (to select an output wavelength and measure the fluorescence intensity).

Figure 5(a) shows a Gaussian fit to the output fluorescence intensity of Oregon Green^{TM}488 excited by different wavelengths of light observed through a 532nm emission filter. This curve matches the expected excitation spectrum of Oregon Green^{TM} as shown in the left curve in Figure 5(b) provided by the vendor [1] and demonstrates the coupling of chromophores to a waveguide.

3.5. Discussion

This section further discusses the mNoC network and qualitatively compares it to ring resonator based networks. First we discuss the scalability of an mNoC by doing theoretical calculation. We follow this by comparing the key device parameters of the two networks.

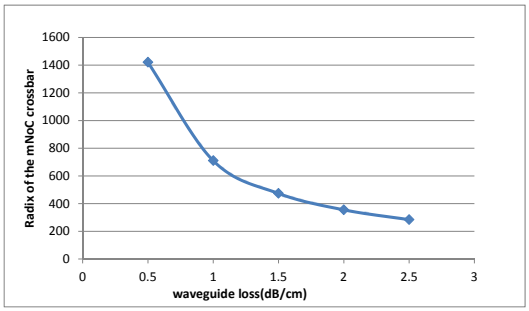

Figure 6: Scalability of mNoC with serpentine layout

3.5.1. Scalability of mNoC As mentioned in the Section 2.2, ring resonator based networks are difficult to scale more than 64 nodes mainly due to the prohibitively high thermal ring power. However, since no rings are used in the mNoC network, an mNoC should scale to a higher radix. In this subsection, we discuss the scalability of an mNoC using theoretical calculations and we show how an mNoc scales up to 256 nodes.

The optical power budget is the difference between the maximum injected power and the extracted power [6]. It determines the maximum insertion loss through the worst case optical path and the number of wavelength channels inside the waveguide. The maximum injected power is limited by nonlinearity of the waveguide itself and the first modulator encountered, while the extracted power is related to the photodetector's sensitivity.

However, for our mNoC technology, QD_LED does not have nonlinearity, and optical output power increases with increasing injection current [33]. We can use all the waveguide power budget by increasing the number of nodes in the system. In order to achieve the maximum optical power budget, we choose to use a photodetector with $0.1\mu W$ sensitivity [26, 37]. Together with $115mW$ maximum injected power budget and 10 wavelengths per waveguide, we get about $50dB$ insertion loss tolerance according to equation (3) in J. Chan et al. [9]. Waveguide's transmission loss is the main source of mNoC's insertion loss, and it is a function of length which roughly defines the number of nodes connected to it. Therefore, a scalability graph can be achieved as shown in Figure 6. From this graph we see that the number of nodes connected by an mNoC increases as the insertion loss decreases. Even with $1.5dB/cm$ loss waveguide, we can easily scale our crossbar up to radix 256. If multilayer silicon integration [6] is used, that number will increase.

3.5.2. Key Parameter Comparison Table 1 summarizes the key parameters for comparison of mNoC to ring resonator based networks (rNoC). From this table we can see the transmitter is at least four times smaller than a ring modulator and the waveguide is half of the pitch size of that used in rNoC.

To estimate the receiver area, we need to know the sensitivity of the photodetector. It determines the minimum input power for the farthest node from the source on the communication path. If the minimum input power can be detected at

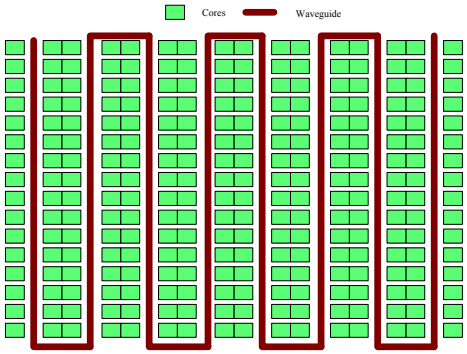

Figure 7: Serpentine layout for 256×256 mNoC

the farthest node, all the other inputs to nodes along the way should also be detected. The sensitivity varies from $80\mu W$ to $0.1\mu W$ [30, 26, 37] for photodetectors used on chip. To match the scalability analysis above, we choose a photodetector with $0.1\mu W$ sensitivity, which equals to about 3 photons with excited state lifetime of $10ps$ and wavelength of $500nm$. If we conservatively assume that the quantum yield of chromophores is 0.1, then 30 photons are absorbed to achieve that sensitivity which requires 30 chromophores. If we assume one chromophore takes $2nm$x$2nm$ area, it is $120nm^2$ in total, which is negligible compared to the waveguide width. For the photodetector itself, the active area can be as small as $0.25\mu m$x$0.25\mu m$ [36, 37]. Compared with ring resonators's area of $10\mu m$x$10\mu m$, the receiver in an mNoC is at least two orders of magnitude smaller.

4. mNoC Architecture

Our proposed chromophore-based technology does not support switching functionality, therefore non-switch-based topologies, such as bus-based crossbars, may be more suitable. We can design an mNoC bus-based crossbar by (i) replacing all the ring resonators in the receiving side with chromophores, (ii) using the photodetector in place of the receivers, and (iii) replacing all the ring modulators with QD_LEDs on the sending side. We can remove the off-chip laser source while keeping the dedicated channels for each of the nodes. This section discusses the detailed design of the mNoC bus.

4.1. Topology.

Single Writer Multiple Reader (SWMR) and Multiple Writer Single Reader (MWSR) are both feasible for an mNoC. However, we choose SWMR over MWSR because no arbitration is required and the design can be simpler. In the SWMR design, to make full use of the big crossbar, we can simply build a flattened crossbar structure. Each core has properly sized buffers to receive packets from other cores. We do not have electrical routers, so the total number of hops is reduced to a single bus traversal in single mNoC systems or fewer overall hops if an mNoC is used in high radix routers.

Ring resonator based crossbars usually use serpentine loop waveguide layout to guide off-chip optical signal on chip and perform broadcast. We can also use the same layout as shown in Figure 7.

4.1.1. Packet Design. mNoC does not require any special design of network packets; however, adding a few bits might benefit power consumption. QD_LEDs use presence of a signal to represent logic 1 and absence of a signal to represent logic 0. The ratio of 1s to 0s affects the power consumption. For applications with more '1's transmitted than '0's, we can invert the representation of signals. Therefore, we can add one invert bit to the packet header to indicate an inverted bit pattern. Other encoding methods can also be explored to reduce the QD_LED's 1-to-0 ratio, but we leave that as future work.

4.1.2. Discussion Compared with ring resonator based crossbar, an mNoC SWMR crossbar changes key NoC parameters and capabilities in the following four ways:

Energy: Overall energy consumption is reduced because the large amount of ring thermal tuning power is removed. Moreover, the off-chip traffic activity-independent laser source is replaced with on-chip QD_LED, where the 1-to-0 emission value ratio and waveguide link utilization play an important role in further reducing the energy consumption.

Area: For mNoC networks, area occupied by optical devices including receivers and transmitters are much smaller than ring resonators. Waveguides are half of the size as those used in rNoC but with about one sixth wavelength density, and therefore become the largest portion of the total area.

For existing 64x64 bus-based crossbars, an mNoC requires 6.4X more waveguides traversing all the nodes once for the same datapath width. These waveguides in mNoCs are half the size of those in ring-based NoCs and on-chip QD_LEDs reduce the two-round or one-round datapath design which doubles the waveguide area. Therefore, an mNoC gives 1.6X area in total compared to a radix-64 ring-based NoC network. However, this area easily fits into a $400mm^2$ chip using the 3D stack proposed by Corona[38].

For a radix-256 mNoC, the total width of the waveguides W_{total} using serpentine layout is shown in Equation 1. W_{pitch} is waveguide pitch, which is $2\mu m$ in our case; N_{nodes} is number of nodes, which is 256, connected by the waveguides; N_p is the number of columns of waveguides which is 8 as shown in Figure 7.

$$W_{total} = W_{pitch} \times \frac{datapath\_size}{wavelength\_density} \times N_{nodes} \times N_{col} \quad (1)$$

If the network datapath size is 256-bit, the waveguide width in total will be about 10 cm which is beyond the width 2cm of the $400mm^2$ die. However, there are three solutions to address this problem. First, if we use multi-layer of nanophotonic fabrication and designs [9, 46], all the waveguides can be easily fit into 4-5 layers.

Second, use a smaller waveguide. A pitch of $200nm$ waveguide will enable us to fit everything on a single layer chip. The waveguide size and transmission loss is a trade-off [39]. With a $200mn$ pitch waveguide, the transmission loss might be

Component	Term	Ring Resonator NoC	mNoC
All	wavelength (nm)	1550	390-750
	thermal tuning power(μW)	(20-100)/ring	0
Transmitter	modulator type	ring modulator	QD_LED
	nonlinearity(mW)	0.6	none
	modulator size(μm^2)	100	21.5(64x64);1.9(256x256)
	modulation rate(GHz)	~ 10	1.1-67
	source efficiency(%)	30	8.9-10.5
	source energy injection rate	1	'1' to '0' ratio
Receiver	filter type	ring resonator	chromophore
	filter size(μm^2)	100	0.625(0.1μW sensitivity)
	sensitivity(μW)	0.1	0.1
Waveguide	#wavelengths	64	10
	pitch(μm)	4	2
	trans. loss(dB/cm)	1-1.5	1-1.5

Table 1: Key Parameters of Ring Resonator NoC and mNoC

bigger than $1dB/cm$. However, using a different layout such as an H-tree can reduce the longest waveguide traverse length from $18cm$ to $4cm$, and therefore reduce the total transmission loss. Designing the $200nm$ waveguide with a satisfying transmission loss is beyond the scope of this paper.

Third, reduce the datapath size of the network. As shown in Equation 1, W_{total} is proportional to $datapath\_size$. Thus, the smaller $datapath\_size$ is, the fewer layers are needed. With a 32-bit datapath, only a single layer integration is enough for all the waveguides. The performance degradation of reduced datapath width is examined in Section 5.

Frequency: The switching frequency of the QD_LED has a range from 1.1 to $67GHz$, as shown in Table 1, depending on the fluorescence lifetime of the selected quantum dots. This frequency decides the fastest switching frequency, and thus data rate, in the network. We leave investigating this potential opportunity for future work.

Broadcast: Every chromophore receiver couples optical power from a waveguide to a photodetector at very low cost ($0.05\mu W$ for a $0.1\mu W$ sensitivity photodetector) regardless of packet types: broadcast, unicast, or multicast. This means an mNoC can support broadcast and multicast with little or no extra cost.

5. Evaluation

This section presents our evaluation of mNoCs using a combination of simulation and analytic methods. We begin by evaluating the performance of mNoCs using synthetic benchmarks in a 256-node system, and compare against alternative topologies implemented in different technologies. Specifically, we compare to a ring resonator based network (rNoC) that uses a 64x64 crossbar and hierarchical clustering to scale to 256 nodes, and to a conventional electrical 256-node 2D mesh. We then evaluate the mNoC when used in a system with MOSI directory protocol using 12 SPLASH benchmarks [41]. This includes performance and energy comparison, datapath impact and longevity discussion.

mNoC	Molecular SWMR crossbar; radix-256; flattened.
rNoC	Ring based SWMR clustered crossbar; radix-64; concentration 4.
eMesh	Electrical mesh structure; 16x16; concentration 1.

Table 2: Networks Evaluated

5.1. Experimental Setup

We create our mNoC topology in the Graphite [28] simulator and run all the simulations in the full simulator mode. The simulation configuration is summarized in Table 3. The total O/E to E/O latency is about 200 ps and is modeled as 1 cycle in the nanophotonic link traversal time. If we assume 256 cores with a die size of $400mm^2$, then the waveguide's total length will be around $18cm$. The speed of light in silicon is about $10cm/ns$, which means $1.8ns$ to travel the longest distance. If the clock rate is $5GHz$, this equals to 9 cycles in the worst case. All the electrical links are modeled as 1 cycle [10] for the alternative networks. The networks we evaluate are listed in Table 2 and each of them has 256 cores in total. The first is our radix-256 mNoC. The second is a radix-64 Firefly [31]-like network—rNoC. We use a normal 4 stages pipeline router for rNoC to connect its four cores in the cluster, and also to connect cores in traditional electrical mesh network.

We model the contention delay of mNoC using the history_tree queue model in Graphite. We instantiate an outgoing and an incoming queue for each node. The outgoing queue is responsible of writing packets onto its dedicated sending channel, while the incoming queue is responsible of reading packets from all the other channels.

5.2. Evaluation with synthetic Benchmarks

We use three synthetic benchmarks: uniform random, bit complement and shuffle traffic from the Graphite simulator to

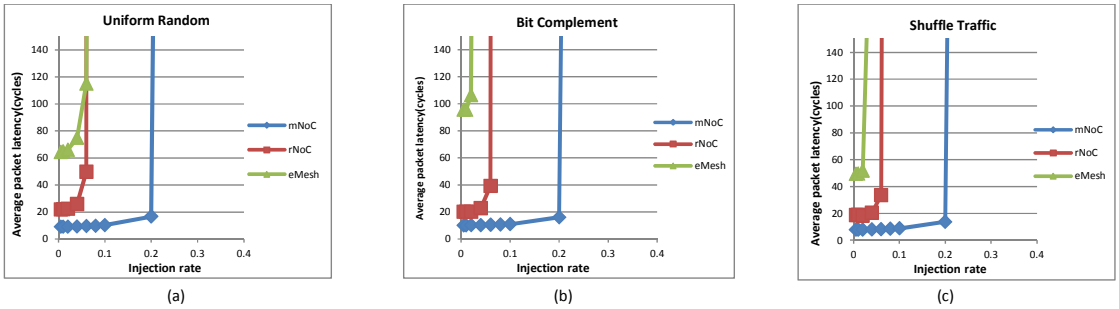

Figure 8: Average Packet Latency for Synthetic traffic pattern (a)Uniform Random, (b)Bit Complement, and (c) Shuffle traffic

Router pipeline stages	4 cycles
Electrical link latency	1 cycles
Optical link latency	1-9 cycles for mNoC; 1-5 for rNoC
Clock	5GHz
Flit size	256-bit
Core model	in-order model, private 32KB L1D, 32KB L1I, 512KB L2 Cache

Table 3: Simulation configuration

QD_LED energy efficiency	10%
QD_LED 1-to-0 ratio	1
Waveguide loss	$1dB/cm$
Coupler loss	$1dB$
Power loss of chromophores	$0.05\mu W$ for $0.1\mu W$ sensitivity
Optical splitter	$0.2dB$

Table 4: Optical energy parameters

perform network only tests. Figure 8(a), (b) and (c) show the average packet latency versus injection rate simulation results for the three 256-node networks. The average latency is measured from when a packet enters the outgoing queue of the source node until it is popped off the incoming queue at the destination node.

In these simulations, all the nodes generate different requests and send them to the corresponding destination. From these three figures we see that when the injection rate increases, the average latency also increases for all three networks. When the network saturates, the average latency dramatically increases due to the queuing delay. These results show that the mNoC has lower average latency compared to the two alternatives. On average, mNoC has half the latency of rNoC and a quarter the latency of eMesh. Furthermore, when the injection rate increases, the mNoC tolerates more traffic. The alternative networks saturate when the injection rate is less 0.08 while the mNoC functions normally until the injection rate reaches 0.2. This is because the mNoC's larger crossbar provides all to all connectivity and deliver packets efficiently.

5.3. Evaluation with SPLASH Benchmarks

As an initial step toward evaluating how mNoC crossbars may influence multi-core computing, we evaluate mNoC's performance and energy against the alternatives with SPLASH workloads in a MOSI directory-based coherence protocol. Our goal is to begin exploring the impact of mNoC as a communication substrate for multicore computing, and to affirm the expected scenario that for workloads that require shared memory accesses mNoCs should perform well. Further discussions of mNoC including impact of datapath width and longevity are also presented.

5.3.1. Performance Comparison From the previous synthetic workloads, we see mNoC's performance advantage under certain traffic patterns: uniform random, bit complement and shuffle traffic. However, in reality, benchmarks might not follow those three patterns. Thus, we also further evaluate mNoC's performance under more realistic workloads. We use 12 multi-threaded benchmarks from SPLASH [41], and run simulations with 256 threads. Figure 9 shows the average packet latency for each of the three networks and Figure 10 shows speedup relative to mesh.

The mNoC crossbar shows significant average latency advantage over the other two alternatives in Figure 9. Its average latency is approximately half of rNoC's and less than one quarter of eMesh's latency. This is because the high radix mNoC crossbar has a flattened structure and it can efficiently transmit packets without taking multiple hops. The improvement of packet latency also gives us better overall system performance. From Figure 10, we see that mNoC performs the best among the networks. On average, it is more than 43% better than eMesh and more than 10% better than rNoC. For some benchmarks such as ocean_c and ocean_nc with relatively high number of shared memory accesses [5], mNoC achieves more than $2X$ better performance compared with eMesh and more than 22% better performance compared with rNoC.

5.3.2. Energy Comparison To show how the new technology affects the energy consumption in an mNoC crossbar, we create an energy model of our crossbar to compare against rNoC. The key parameters used in the model are listed in Table 4. To match the clock rate, we evaluate the quantum dots with lifetime of $200ps$ [19, 4] in our model which corresponds to signal switching frequency of $5GHz$ in the network. We assume the source energy efficiency is 10% [2].

A notable feature of mNoC is that the QD_LED light source

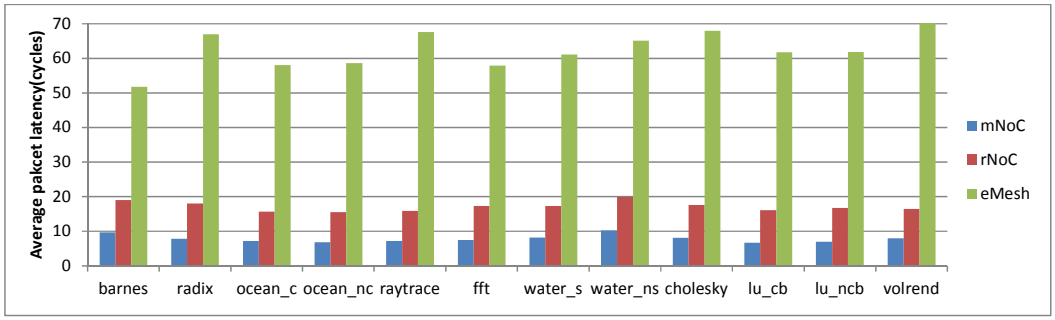

Figure 9: Average Packet Latency of SPLASH Benchmarks

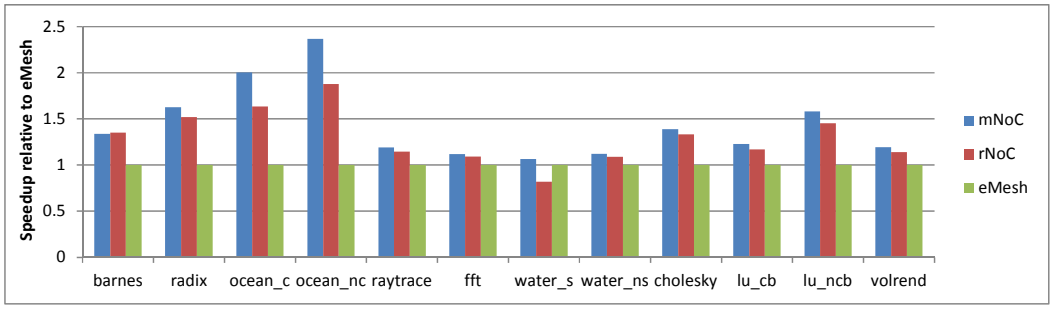

Figure 10: Speedup Comparison of SPLASH Benchmarks

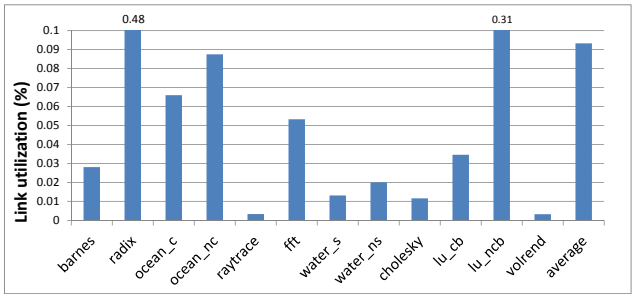

Figure 11: Link utilization of SPLASH Benchmarks

in an mNoC only consumes optical energy when the optical link is in use. No static energy is consumed if there is no traffic transmitted on the link. This means the link utilization is an important factor for light source power. With 100% link utilization in serpentin layout, the peak power of a radix-256 mNoC is about 130W, which is much smaller comapred to the 225W peak power of an 8x8 eMesh with 50% input load [13]. The peak can be greatly reduced by using a different waveguide layout, such as an H-tree. Furthermore, link utilization is not always 100%. From running all the 12 SPLASH benchmarks, we gathered the link utilization of mNoC as shown in Figure 11. On average the link utilization is 0.1%, which is low due to limited number of shared memory accesses [5], abundant on-chip link resources (e.g., caches) and wide datapath and flit size (both are 256-bit).

Given the link utilization, we conservatively calculate the QD_LED source energy by multiplying by an additional factor of 10 to account for chromophore photoreceivers' loss due to inefficient RET and noise margin to favor the alternative networks. Additionally, the 1-to-0 ratio of QD_LED can also affect the source energy consumption, because only the logic '1' values actually consume energy. We assume the ratio to be 1 which means the worst case without any bit inverse scheme. For ring resonator based devices, all the optical parameters are similar to Flexishare [30]. For electrical circuit and links, we refer to models presented elsewhere [30, 31, 14]. Furthermore, for rNoC we optimistically assume that there are splitters before ring receivers to get the right amount of power to reduce influence of insertion loss caused by serial ring resonators. We add one more network—clustered_mNoC(c_mNoC) for comparison. c_mNoC has the same organization as r_NoC, but instead of using ring resonator-based technology, we use the proposed molecular scale technology.

Figure 12 shows the average energy results across all 12 SPLASH benchmarks. From Figure 12(a), we can see that the average energy of a radix-256 mNoC crossbar is reduced by 49% compared to the radix-64 clustered rNoC (256 total cores). This improvement is due to removal of ring resonator thermal tuning as shown in the energy breakdown graph Figure 12(b) and (c), and the QD_LED dependence on traffic. mNoC couples power into the waveguide when needed as opposed to constantly coupling power from an off-chip laser source. A radix-64 clustered mNoC network (C_mNoC) as shown in Figure 12(a) uses only 23% of the energy as rNoC, which has a similar structure but uses different techonology.

5.3.3. Datapath Width Impact A wide datapath is efficient for utilizing on-chip bandwidth, however, it also increases

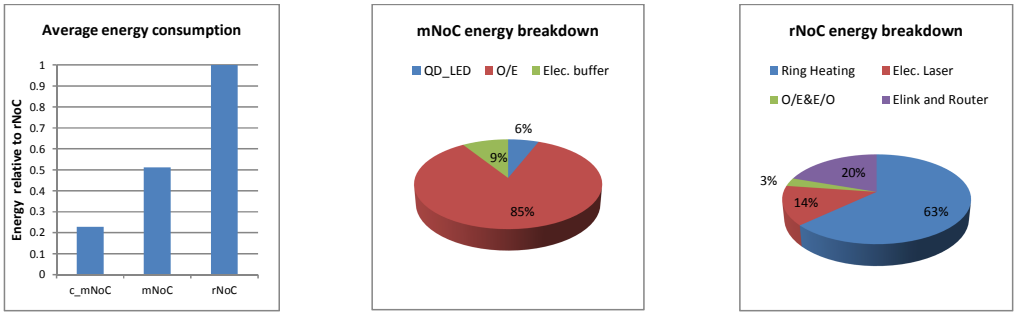

Figure 12: Energy comparison (a) average energy consumption across all SPASH benchmarks, (b) radix-256 mNoC , and (c) radix-64 clustered rNoC average energy breakdown.

cost such as area as analyzed in Section 4.1.2. We compare the performance and energy of SPLASH benchmarks with different NoC datapath widths and the results are shown in Figure 13. Figure 13(a) indicates the performance decreases for smaller datapath widths, with an average 5% performance degradation for width 32. The performance decrease is due to the increased serilization latency of a smaller datapath. Figure 13(b) shows the smaller datapath width can efficiently reduce energy by up to 77% for width 32.

5.3.4. Longevity Despite the small size and low power consumption of chromophores, they also have a photobleanching problem [20] which might reduce the longevity of the mNoC network. However, we can overcome this problem (1) by increasing the total number of chromophores integrated and reducing the intensity of QD_LED source power to still match the desired sensitivity of the photodetector, and (2) using encapsulation [29, 8] .

A single chromophore takes $4nm^2$ area and we allocate as many chromophores as possible into the area $0.625\mu m^2$ of a photodetector. If we assume 5 layers of chromophores can be stacked together which gives us a thickness of $10nm$, the total number of chromophores $N_{chromophores}$ is 0.78 million. High performance chromophores can be excited more than $10^8(L_{single})$ times before completely photobleaching [23, 22, 24]. If we consider the chromophores utilization and also use the encapsulation to extend the longevity by a factor of 10 (F_{encap}) [8], we will get about 10 years longevity of the chromophore receivers according to equation 2. In the equation, N_{year_cycles} is total cycles with a $5GHz$ clock rate for one year, U_{link} is link utilization and $P_{1\_to\_0}$ is the percentage of signal '1's(50%) with a 1_to_0 ratio of 1.

$$N_{years} = \frac{N_{chromophores} \times L_{single}}{N_{year\_cycles} \times U_{link} \times P_{1\_to\_0}} \times F_{encap} \qquad (2)$$

6. Architectural Implications

In power aspect, molecular optical technology enables us to reduce the on-chip energy 77% by keeping the same radix of the crossbar, or scale to 256x256 SWMR crossbar with 49% less energy consumption. In performance aspect, the average latency is significantly shortened when a big crossbar is used. These properties of mNoCs can have significant implications on router micro architecture and system architecture design. In this section, we discuss some of the implications.

6.1. High radix router

The high radix router is the trend for future network-on-chips [15]. It greatly reduces the number of hops in the network and provides better latency at lower cost. After optics have been introduced to Network-on-chip, they start to play an more and more important role in high radix router designs [7].

Molecular technology introduced by this paper is a good candidate to build a high radix router. The 256x256 big crossbar enables us to provide high all-to-all connectivity within a router while keeping the energy consumption within a low budget. With a few such nanophotonic crossbars it is possible to construct higher radix routers, or for systems with <=256 nodes to utilize a single crossbar. Furthermore, with the development of on-chip optics, waveguides with lower transmission loss will allow us to scale to even higher radix crossbar as shown in Figure 6.

6.2. Cache Coherence

One feature of our mNoC is that it broadcasts packets at no extra cost. Therefore, broadcast or multicast can be efficiently supported with an mNoC without any significant change in energy consumption. Some simulation results from Section 5.3.1 show the average latency benefit and performance speedup when utilizing MOSI directory protocol. Positive results are obtained even for workloads without much shared memory accesses which communicate data. This opens the possibility to efficiently support broadcast based coherence protocols and simplify the cache coherence design. However, further research is required to explore which protocol makes best use of mNoC features. Developing a customized cache protocol to mNoC is part of our ongoing work.

7. Related Work

In this section, we summarize different nanophotonic NoC designs and how they address the laser energy consumption

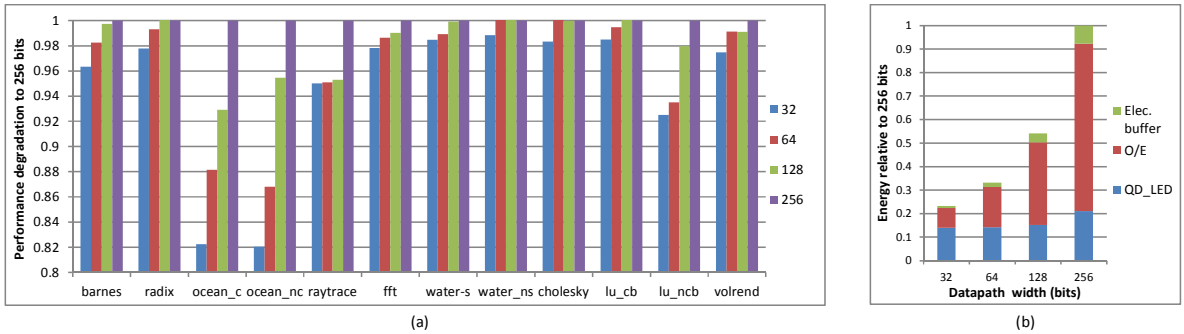

Figure 13: Performance (a) and average energy across all SPLASH benchmarks (b) comparison of different datapath widths

and ring thermal tuning problem while maintaining good performance.

Among many topologies used in nanophotonic NoC, the bus-based crossbar is the most popular one because of its simplicity. There are basically three structures for bus-based crossbars. One is Multi-Write-Single-Read structure proposed by Corona [38]. Each node has its own channel to read which all the other channels can write into, therefore global write arbitration is required. A token-based arbitration method is proposed where the token represents the right to modulate the node's data to deliver. However, when the contention is low, a node can wait for a long time to obtain a token and therefore it's not efficient. As we mentioned before, Corona uses $1024K$ ring resonators and the thermal tuning energy is high. Similar to Corona, Binkert et al. [7] also use MWSR structure.

A Single-Write-Multi-Read shared-bus crossbar is proposed by Kirman et al. [18] and Pan et al. [31]. All the nodes send signals into their own channel and other nodes listen to it and receive them. To avoid high power consumption from coupling power from the laser of the sending channel to all the receiving channels (broadcast), Pan et al. add a reservation-assisted SWMR bus to reduce power consumption by avoiding delivering power to untargeted channels. In this way, the broadcast becomes unicast and laser power is saved.

However, both MWSR and SWMR need one channel for each node and the same amount of ring resonators for each channel to modulate and filter out signals. Multi-Write-Multi-Read structure proposed in Flexishare [30] combines both and proposes a reduced number of channels design which uses a token-based mechanism for arbitration on the sending side and reservation channel on the receiving side. As a result, both the laser loss and ring heating energy are reduced. However, they still take more than 50% of total energy in most cases and the router design is more complex than both MWSR and SWMR with a higher power consumption.

Some other topologies besides shared-bus crossbar are also proposed based on ring resonator switches. For example, Phastlane [12] presents a 2D grid NoC of optical crossbar switches. The switch uses optical-level, source-based switch control supported by an electrical network to reduce the latency. THOE [43] is a torus-based hierarchical hybrid NoC. It employs some new techniques such as floorplan optimization, an adaptive power control mechanism and hybrid routers with a low-power optical switching fabric. Since these ring resonator switches-based designs do not use shared channels, the number of ring resonators attaching to channels is greatly reduced. However, they usually have their own problems. For examples, in Phastlane the output port arbitration and packets buffering are performed electrically, and if there is not enough buffering space available, packets will be dropped and performance will be hurt. In addition to that, the time of driving resonators constitute a big part of delay in the critical path through the Phastlane's router. In THOE, instead of off-chip CW laser source, it uses on-chip VCSELs, which emits light vertically and requires integrated mirrors and complicated lithographic technologies to transfer light to the horizontal surface [18]. Each VCSEL pill has a diameter of $55\mu m$. Therefore, the fabrication cost and scalability might be a problem.

From the above related work, we can see that no matter what design/topology is proposed, a lot of effort is made to reduce the laser power and the number of ring resonators in current nanophotonic NoC because those two parts take a major portion of the total energy consumption. However, for crossbar topologies, laser power and ring thermal tuning still take about half of the total power consumption. For non-crossbar topologies, other problems exist such as fabrication cost. The scalability and optical computing problems remain open challenges.

8. Conclusion and Future Work

Current ring resonator based nanophotonics NoCs provide for higher bandwidth and more power efficient designs compared with traditional CMOS NoCs. They often use an external laser source, ring resonators for signal modulation and filtering, and waveguides for transmission. However, they still suffer from high static power consumption due to losses in the external laser source and thermal ring tuning, and limited network scalability.

In this paper, we propose to use emerging molecular scale devices to construct nanophotonic networks— Molecular-scale Network-on-Chips (mNoCs). The molecular scale devices include quantum dot LEDs, which provide electrical to

optical signal modulation, and chromophores, which provide optical signal filtering for receivers. The chromophores replace the ring resonators and the quantum dot LED replaces the external laser source used in current nanophotonic NoCs. We present a Single Writer Multiple Reader (SWMR) bus based crossbar mNoC and show that without limitations of current nanophotonic networks, an mNoC crossbar can easily scale to a radix-256 crossbar.

We evaluate mNoC with both synthetic benchmarks and real workloads SPLASH benchmarks. Synthetic benchmarks simulation results show that compared with two alternative networks rNoC and eMesh, mNoC has greatly reduced average latency cycles(half at least) and has higher tolerance for network traffic. mNoC also exhibits 10% speedup improvement over rNoC and 43% speedup improvement over eMesh across 12 SPLASH benchmarks on average. Furthremore, mNoC trades static energy for dynamic enegy and greatly reduces energy consumption. An mNoC can achieve 77% reduction in energy for a 64x64 crossbar compared to similar ring resonator based designs. Additionally, an mNoC can scale to a 256x256 crossbar with 49% of energy reduction. A large single crossbar allows for the possibility of high radix routers and efficient broadcast based directory protocols. We also discuss implications of the new mNoC crossbar on overall system design.

mNoC presents several areas for future investigations, here we present three examples. First, evaluate different cache coherence protocols with mNoC, such as Token protocol vs AMD hammer, modifying the protocols if necessary to better utilize mNoC's properties (e.g., broadcast/multicast). Second, design a high radix router with mNoC and evaluate its performance, power, and area etc. Third, augment mNoC with nanophotonic computational abilities to further reduce energy consumption or improve performance.

References

[1] "http://www.invitrogen.com/site/us/en/home/support/product-technical-resources/product-spectra.12650lip.html," Life Technologies Corporation.

[2] "http://www.qdvision.com/qled-technology," QD Vision, Inc.

[3] P. Anikeeva *et al.*, "Electronic and excitonic processes in light-emitting devices based on organic materials and colloidal quantum dots," *Physical Review B*, vol. 78, no. 8, p. 085434, 2008.

[4] R. Arians *et al.*, "Electrically driven single quantum dot emitter operating at room temperature," *Applied Physics Letters*, vol. 93, no. 17, pp. 173 506–173 506, 2008.

[5] N. Barrow-Williams, C. Fensch, and S. Moore, "A communication characterisation of splash-2 and parsec," in *Workload Characterization, 2009. IISWC 2009. IEEE International Symposium on*. IEEE, 2009, pp. 86–97.

[6] A. Biberman *et al.*, "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.

[7] N. Binkert *et al.*, "The role of optics in future high radix switch design," in *Proceeding of the 38th annual international symposium on Computer architecture*. ACM, 2011, pp. 437–448.

[8] G. Canton *et al.*, "Modified stöber synthesis of highly luminescent dye-doped silica nanoparticles," *Journal of Nanoparticle Research*, vol. 13, no. 9, pp. 4349–4356, 2011.

[9] J. Chan *et al.*, "Physical-layer modeling and system-level design of chip-scale photonic interconnection networks," *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on*, vol. 30, no. 10, pp. 1507–1520, 2011.

[10] G. Chen *et al.*, "Predictions of cmos compatible on-chip optical interconnect," *Integration, the VLSI journal*, vol. 40, no. 4, pp. 434–446, 2007.

[11] S. Chen *et al.*, "Bistability and self-pulsation phenomena in silicon microring resonators based on nonlinear optical effects," *Optics Express*, vol. 20, no. 7, pp. 7454–7468, 2012.

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

[13] Y. Jin, E. Kim, and K. Yum, "Peak power control for a qos capable on-chip network," in *Parallel Processing, 2005. ICPP 2005. International Conference on*. IEEE, 2005, pp. 585–592.

[14] A. Joshi *et al.*, "Silicon-photonic clos networks for global on-chip communication," in *Proceedings of the 2009 3rd ACM/IEEE International Symposium on Networks-on-Chip*. IEEE Computer Society, 2009, pp. 124–133.

[15] J. Kim *et al.*, "Microarchitecture of a high-radix router," in *ACM SIGARCH Computer Architecture News*, vol. 33, no. 2. IEEE Computer Society, 2005, pp. 420–431.

[16] K. Kim *et al.*, "Shell layer dependence of photoblinking in cdse/znse/zns quantum dots," *Applied Physics Letters*, vol. 98, p. 012109, 2011.

[17] L. Kim *et al.*, "Contact printing of quantum dot light-emitting devices," *Nano letters*, vol. 8, no. 12, pp. 4513–4517, 2008.

[18] N. Kirman *et al.*, "Leveraging optical technology in future bus-based chip multiprocessors," in *Proceedings of the 39th Annual IEEE/ACM International Symposium on Microarchitecture*. IEEE Computer Society, 2006, pp. 492–503.

[19] T. Kümmell *et al.*, "Electrically driven room temperature operation of a single quantum dot emitter," in *Proceedings of SPIE*, vol. 7211, 2009, p. 72110G.

[20] J. Lakowicz, *Principles of fluorescence spectroscopy*. Springer, 2006, vol. 1.

[21] J. Lakowicz and B. Masters, "Principles of fluorescence spectroscopy," *Journal of Biomedical Optics*, vol. 13, p. 029901, 2008.

[22] H. Langhais, "Cyclic carboxylic imide structures as structure elements of high stability novel developments in perylene dye chemistry," *Heterocycles-Sendai Institute of Heterocyclic Chemistry*, vol. 40, no. 1, p. 477, 1995.

[23] H. Langhals, "Control of the interactions in multichromophores: Novel concepts. perylene bis-imides as components for larger functional units," *Helvetica chimica acta*, vol. 88, no. 6, pp. 1309–1343, 2005.

[24] H. Langhals, J. Karolin, and L. Johansson, "Spectroscopic properties of new and convenient standards for measuring fluorescence quantum yields," *Journal of the Chemical Society, Faraday Transactions*, vol. 94, no. 19, pp. 2919–2922, 1998.

[25] Z. Li *et al.*, "Spectrum: A hybrid nanophotonic electric on-chip network," in *Design Automation Conference, 2009. DAC'09. 46th ACM/IEEE*. IEEE, 2009, pp. 575–580.

[26] Z. Li *et al.*, "Iris: A hybrid nanophotonic network design for high-performance and low-power on-chip communication," *ACM Journal on Emerging Technologies in Computing Systems (JETC)*, vol. 7, no. 2, p. 8, 2011.

[27] S. Liao *et al.*, "36 ghz submicron silicon waveguide germanium photodetector," *Optics Express*, vol. 19, no. 11, pp. 10 967–10 972, 2011.

[28] J. Miller *et al.*, "Graphite: A distributed parallel simulator for multicores," in *High Performance Computer Architecture (HPCA), 2010 IEEE 16th International Symposium on*. IEEE, 2010, pp. 1–12.

[29] H. Ow *et al.*, "Bright and stable core-shell fluorescent silica nanoparticles," *Nano letters*, vol. 5, no. 1, pp. 113–117, 2005.

[30] Y. Pan, J. Kim, and G. Memik, "Flexishare: Channel sharing for an energy-efficient nanophotonic crossbar," in *High Performance Computer Architecture (HPCA), 2010 IEEE 16th International Symposium on*. IEEE, 2010, pp. 1–12.

[31] Y. Pan *et al.*, "Firefly: illuminating future network-on-chip with nanophotonics," *ACM SIGARCH Computer Architecture News*, vol. 37, no. 3, pp. 429–440, 2009.

[32] H. Park *et al.*, "A hybrid algainas-silicon evanescent preamplifier and photodetector," *Opt. Express*, vol. 15, no. 21, pp. 230–232, 2007.

[33] I. Park *et al.*, "Green light-emitting diodes with self-assembled in-rich ingan quantum dots," *Applied Physics Letters*, vol. 91, p. 133105, 2007.

[34] G. Priem *et al.*, "Nonlinear effects in ultrasmall silicon-on-insulator ring resonators," vol. 6183, Apr. 2006.

[35] K. Rurack and M. Spieles, "Fluorescence quantum yields of a series of red and near-infrared dyes emitting at 600- 1000 nm," *Analytical chemistry*, vol. 83, no. 4, pp. 1232–1242, 2011.

[36] S. Sahni *et al.*, "Germanium-on-soi photo-detector based on an fet structure," in *CLEO*. Optical Society of America, 2007.

[37] S. Sahni *et al.*, "Junction field-effect-transistor-based germanium photodetector on silicon-on-insulator," *Opt. Lett.*, vol. 33, pp. 1138–1140, May 2008.

[38] D. Vantrease *et al.*, "Corona: System implications of emerging nanophotonic technology," *ACM SIGARCH Computer Architecture News*, vol. 36, no. 3, pp. 153–164, 2008.

[39] C. Wang and L. Lin, "Nanoscale waveguiding methods," *Nanoscale Research Letters*, vol. 2, no. 5, pp. 219–229, 2007.

[40] J. Wang *et al.*, "Evanescent-coupled ge pin photodetectors on si-waveguide with seg-ge and comparative study of lateral and vertical pin configurations," *Electron Device Letters, IEEE*, vol. 29, no. 5, pp. 445–448, 2008.

[41] S. Woo *et al.*, "The splash-2 programs: Characterization and methodological considerations," in *ACM SIGARCH Computer Architecture News*, vol. 23, no. 2. ACM, 1995, pp. 24–36.

[42] V. Wood and V. Bulović, "Colloidal quantum dot light-emitting devices," *Nano Reviews*, vol. 1, no. 0, 2010.

[43] Y. Ye *et al.*, "A torus-based hierarchical optical-electronic network-on-chip for multiprocessor system-on-chip," *ACM Journal on Emerging Technologies in Computing Systems (JETC)*, vol. 8, no. 1, p. 5, 2012.

[44] T. Yin *et al.*, "40gb/s ge-on-soi waveguide photodetectors by selective ge growth," in *Optical Fiber communication/National Fiber Optic Engineers Conference, 2008. OFC/NFOEC 2008. Conference on*. IEEE, 2008, pp. 1–3.

[45] P. Yupapin, C. Teeka, and P. Chitsakul, "Mathematical simulation of nonlinear effects in micro ring resonator," in *Emerging Technologies-Nanoelectronics, 2006 IEEE Conference on*. IEEE, 2006, pp. 316–321.

[46] X. Zhang and A. Louri, "A multilayer nanophotonic interconnection network for on-chip many-core communications," in *Design Automation Conference (DAC), 2010 47th ACM/IEEE*. IEEE, 2010, pp. 156–161.