

mNoC: Large Nanophotonic Network-on-Chip Crossbars with Molecular Scale Devices¹

JUN PANG, Duke University
CHRISTOPHER DWYER, Duke University
ALVIN R. LEBECK, Duke University

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, and waveguides. 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 — Molecular-scale Network-on-Chip (mNoC). We leverage on-chip emitters such as 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 88% reduction in energy for a 64×64 crossbar compared to similar ring resonator based designs. Additionally, an mNoC can scale to a 256×256 crossbar with an average 10% performance improvement and 54% energy reduction.

Categories and Subject Descriptors: C.1.2 [Multiple Data Stream Architectures (Multiprocessors)]: Interconnection architectures (e.g., common bus, multiport memory, crossbar switch)

General Terms: Design, Experimentation, Performance

Additional Key Words and Phrases: Nanophotonic, network-on-chip, crossbar, chromophores, molecular-scale devices

ACM Reference Format:

Jun Pang, Christopher Dwyer, and Alvin R. Lebeck, 2014. mNoC: Large Nanophotonic Network-on-Chip Crossbars with Molecular Scale Devices. *ACM J. Emerg. Technol. Comput. Syst* V, N, Article A (January YYYY), 25 pages.

DOI: <http://dx.doi.org/10.1145/0000000.0000000>

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

¹This paper was an extension of "Exploiting Emerging Technologies for Nanoscale Photonic Networks-on-Chip", published in Sixth International Workshop on Network on Chip Architectures (NoCArc-13), December 2013.

Author's addresses: Jun Pang and Alvin R. Lebeck, Department of Computer Science, Duke University; Christopher Dwyer, Department of Electrical Engineering, Duke University.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© YYYY ACM 1539-9087/YYYY/01-ARTA \$15.00

DOI: <http://dx.doi.org/10.1145/0000000.0000000>

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 64×64 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 $N \times N$ crossbar connectivity ($N \leq 64$), 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) [Pang et al. 2013]. The specific molecular-scale devices we utilize are quantum dot LEDs (QD_LED) and chromophores. Our analysis in the paper is for QD_LED, however, our methods can be applied to any on-chip emissive light source (e.g., QD_LED, VCSEL, and thin-film edge emitting laser etc.). 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. We choose this structure for mNoC not only because it provides full connection between all the input and output ports with low latency, but also because the main component chromophores do not have switching property and bus-based crossbar are more suitable. Our evaluation shows that an mNoC can achieve more than 88% reduction in energy for a 64×64 crossbar compared to ring resonator based designs. Additionally, for 54% less energy budget than a ring resonator 64×64 crossbar, an mNoC can scale to a 256×256 crossbar. For other topologies, such as Multiple Writer Multiple Reader (MWMR) and Multiple Writer Single Reader (MWSR), we present in our prior work [Pang et al. 2013] and due to page limits, we do not include the discussion here.

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.

The remainder of this paper is organized as follows. Section 2 reviews ring resonator NoCs and the motivation of 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 [Miller et al. 2010]. Section 6 discusses the architectural implications of mNoC on router microarchitecture and system architecture. We conclude in Section 7.




Fig. 1. Main Optical Components of Ring Resonator Nanophotonic Network-on-chip

2. RING RESONATOR-BASED NANOPHOTONIC NETWORKS-ON-CHIP

Ring resonator based NoCs are the predominate approach used in recent studies of nanophotonic on-chip networks [Pan et al. 2010; Joshi et al. 2009; Pan et al. 2009; Vantrease et al. 2008]. 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 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. Ring resonators require non-negligible ring trimming power [Pan et al. 2010; Nitta et al. 2011; Xu et al. 2012].

Compared with all electrical NoC designs, most current nanophotonic NoCs exhibit better performance with reduced power consumption [Vantrease et al. 2008; Pan et al. 2009; Kirman et al. 2006]. 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 ring trimming power. Ring resonators have different resonant wavelengths due to variations in temperature. This problem can be addressed by using thermal tuning and/or current injection, which are techniques called trimming. The ring trimming power is assumed to be fixed for a single ring — about $20\mu W$ for a $20K$ temperature range [Joshi et al. 2009; Vantrease et al. 2008; Biberman et al. 2011]. However according to previous research [Nitta et al. 2011], the total trimming power has a nonlinear relationship with ring count and it is related to the total die area. For a $20K$ temperature range, the maximum trimming power for a 64-bit radix-64 256-core system with $484mm^2$ die area is $103W$ ($51.6W$ on average). Even for $400mm^2$ die area, the maximum heating is $98.9W$ ($49.4W$ on average). When the radix increases with more cores integrated on chip, the ring trimming power will become even higher. This high trimming 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 significant portion to the total power. According to analysis from Pan et al. [Pan et al. 2010], 36% of the total energy is consumed in the electrical laser source for a radix-32 conventional nanophotonic crossbar. This is not only because of 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 (i.e., caches), the application characteristics and the high speed of the optical communication structure 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 [Pan et al. 2010]: two-round data channels allow the sender to modulate in the first round and the receiver 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 [Li et al. 2009; Chan et al. 2011; Biberman et al. 2011] and ring resonators [Chen et al. 2012; Priem et al. 2006; Yupapin et al. 2006], 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 [Biberman et al. 2011], while the maximum injected waveguide power without causing nonlinearity is $115mW$ [Biberman et al. 2011].

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.

2.3. Related Work

In this section, we summarize different nanophotonic NoC designs and how they address the laser energy consumption 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 [Vantrease et al. 2008]. 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. [Binkert et al. 2011b] also use MWSR structure.

A Single-Write-Multi-Read shared-bus crossbar is proposed by Kirman et al. [Kirman et al. 2006] and Pan et al. [Pan et al. 2009]. 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 [Pan et al. 2010] 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 [Cianchetti et al. 2009] 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 [Ye et al. 2012] 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. Le Beux et al [Le Beux et al. 2013; Le Beux et al. 2010] propose optical switch based λ -router network and explore how to exploit some regular properties of switch-based ring resonators optical NoC to build 3D architecture to reduce the on-chip resources such as routing elements, and laser sources etc. Since these ring resonator switch-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 example, 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 [Kirman et al. 2006]. Each VCSEL pill has a diameter of $55\mu m$. Therefore, the fabrication cost and scalability might be a problem. In Le Beux et al 's work, it is unclear how power consumption scales when the network scales.

From the above related work, we note that no matter what design/topology is proposed, a lot of effort is made to reduce the static power [Ramini et al. 2014], such as laser power and thermal tuning in current nanophotonic NoC because it 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.

There are other optical interconnection networks which do not use ring resonators as the basic building elements. For example, Ladouceur et al [Liboiron-Ladouceur et al. 2008; Liboiron-Ladouceur et al. 2011] use semiconductor optical amplifier (SOA) as a basic switching element to build scalable optical networks. However, it is mainly used to connect off-chip components and it is hard to compare with on-chip interconnection networks.

The remainder of this paper explores one potential alternative to ring resonator-based nanophotonic NoC to address its limitations.

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 on-chip emitter such as quantum dot LEDs, and chromophores. The quantum dot LEDs provide electrical to optical signal modulation and the chromophores 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 along the communication path. As shown in this figure, when Core 2 sends a packet to Core N, the packet first goes through the modulator driver array to get properly modulated electrical signal to drive Quantum Dot LEDs (QD_LED). QD_LEDs then convert the electrical signal to optical signal. They provide both light source and modulation in a single device. All QD_LEDs from Core 2 inject light into the waveguide through a coupler. Through the SWMR broadcast waveguide, the packet is broadcast to all the other cores. At the receiving side of each core, chromophores deposited on the waveguide work as filters and couple to a photodetector for O/E conversion and the electrical signals are amplified through electrical receiver array. More electrical circuits are required after the receiver array (not shown here), for example, the destination address of the packet is compared with the receiving cores' ID. If they do not match, this packet will be discarded. Or else it will be saved into buffers. In our example here, Core N will save the packet into its buffers. In this paper, we use commercially available chromophores that operate in the wavelength range (400-800nm) 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 5 instead of 64 commonly used in current nanophotonic NoCs.

3.1. Transmitter: Quantum-Dot-LED

Transmitters are composed of silicon compatible QD_LEDs [Gopal et al. 2009] which inject light directly into the waveguide and provide both the light source and modulation in a single on-chip device. QD_LEDs operate as a current controlled light source, more current leads to more photons. We choose QD_LEDs over other comparable technology such as VCSEL, because VCSEL is less power efficient and it requires integrated mirrors and complicated lithographic technologies, thus increases integration and scalability difficulty [Kirman et al. 2006]. In contrast to a continuous off-chip laser

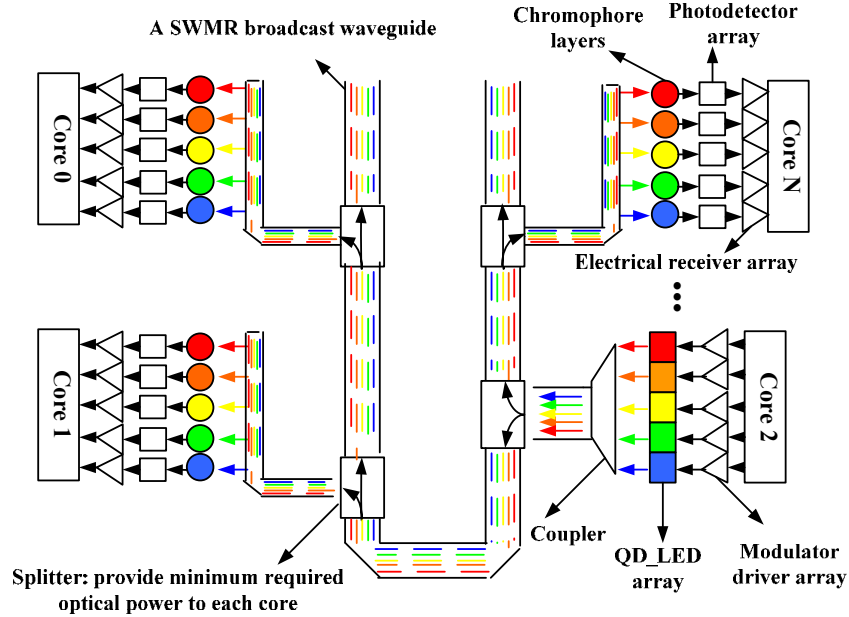


Fig. 2. mNoC Structure and Main Components of a SWMR Broadcast Waveguide for an N-Core System (each core has such a waveguide, not shown for core 0, 1, 3-N)

source, an on-chip QD_LED consumes power only when the optical link is utilized (i.e., transmitting logic '1' but not '0').

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, good stability (the output signal is resistant to degradation caused by effects such as photobleaching) [Wood and Bulović 2010] and fast excitation rate [Wood and Bulović 2010; Hargart et al. 2013]. Important properties to consider when using QD_LED in mNoCs include: energy injection rate, external quantum efficiency, modulation rate, and size.

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.

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 is a very important parameter for calculating the wall-plug efficiency of QD_LED as a light source. The wall-plug efficiency is determined by the external quantum efficiency (EQE) and voltage drop across the device. Mashford et al. [Mashford et al. 2013] demonstrate QD_LEDs with 18% EQE and 2V voltage drop, which suggests a nearly 18% wall-plug efficiency. To calculate the dynamic power for modulation, we use equation $P_D = \alpha CV^2 f$, where α is the activity factor. α is a product of the signal-1 ratio (the percentage of transmitted logic '1' signal among all transmitted '1' and '0')

logic signals) and actual link-utilization of the crossbar. The actual link utilization in SWMR mNoC for the SPLASH benchmarks is as low as 0.1% and the signal-1 ratio is conservatively assumed to be 0.5 (half logic '1' and half logic '0' signal). The dynamic switching power will not become a dominant factor.

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 [Kim et al. 2011] and varies from $15ps$ to $900ps$ for different shells [Kümmell et al. 2009; Arians et al. 2008]. The corresponding modulation rate is in the range $1.1GHz$ - $67GHz$. However, the fastest one which has been demonstrated so far is $2GHz$ from Hargart et al. [Hargart et al. 2013]. In our model, we use $5GHz$ to match clock rates of other NoC work. The optical emission power from QD_LED at this rate depends on the loss path and the photodetector's sensitivity. We can always integrate more Quantum dots to provide more power. In a few milliwatts output range, QD_LED does not exhibit nonlinearity [Park et al. 2007b].

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 $100nm$ [Anikeeva et al. 2008]. By controlling the effective area of QD_LED, we can achieve different amounts of output power as needed. For example, if we want to drive 255 photoreceivers with sensitivity of $1\mu W$ along a $18cm$ long waveguide ($1dB/cm$), the output power from the QD_LED should be $3.45mW$. With 0.8 internal quantum efficiency [Anikeeva et al. 2008], that equals to 2 million photons with wavelength $500nm$ and decay time $200ps$. To estimate the area of QD_LEDs for this $3.45mW$ output power, we can use contact printing fabrication [Kim et al. 2008]. Ten 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 2 million QDs, that is about $215\mu m^2$. However, if we have a $10dB$ insertion loss path which is roughly what the Corona 64×64 crossbar has [Biberman et al. 2011] and drive 63 photoreceivers, then the area becomes $19\mu m^2$. Compared with ring resonator with diameter in the μm range, our area is in the same order of magnitude.

3.2. Receiver: chromophores

An mNoC utilizes chromophores to filter out optical signals and couple to photodetectors to perform O/E conversion. For an mNoC, signal transmission from the waveguide to the chromophores and from the chromophores to the photodetector are both through near field evanescent coupling [Wang et al. 2008; Park et al. 2007a; Liao et al. 2011]. The chromophore couplings take near instant time, and thus the delay can be ignored.

Chromophores have different properties from ring resonators. First, they are very small ($\sim 1nm$) as opposed to μm range radius for a ring resonator, so there is an area range 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 trimming power and their energy loss is proportional to the input energy.

An optical signal transmits from chromophores to waveguides and then to photodetectors through near field evanescent coupling. To prevent light from directly coupling between the waveguide and the photodetector, the photodetector must operate at a longer wavelength outside the visible light range, such as the near-infrared range [Rurack and Spieles 2011]. We need two layers of chromophores as shown in Figure 3. The first layer works as filters to select the desired visible light from the waveguide (99% absorption probability according to our Matlab simulation). The second layer works as a wavelength converter to red-shift the selected visible wavelength to near-infrared through resonance energy transfer (RET, over 95% efficiency), the near field transfer




Fig. 3. mNoC Layered Chromophore Receiver Design

of energy between compatible chromophores [Valeur 2001]. Ge photodetectors as those used in [Pan et al. 2009; Pan et al. 2010] have very high coupling efficiency (0.1dB) in the near-infrared and thus a good candidate. Since RET happens much faster ($9ps$) than the clock period ($200ps$), the delay can be ignored;

3.3. Waveguide

The current photonic NoC uses silicon waveguide [Joshi et al. 2009] with a pitch of $4\mu m$ and transmission loss of $0.2-1.5dB/cm$. Since silicon is opaque to visible light, we use a subwavelength-diameter silica (SiO_2) waveguide [Tong et al. 2003] instead. Its transmission loss is $0.05-1dB/cm$ when the diameter is $450-950nm$. Depending on the desired crosstalk due to chromophores excitation spectrum overlap, we can transmit 8 or 9 wavelengths with $-5dB$ crosstalk², 7 wavelengths with $-7dB$ crosstalk, or 5 wavelength with $-10dB$ crosstalk for the neighboring channels. The crosstalk can be further reduced by using simple thresholding circuits after the photoreceivers. While this introduces more waveguides, several techniques exist to accommodate these additional waveguides (e.g., 3D stack [Vantrease et al. 2008], multi-layer nanophotonic fabrication [Biberman et al. 2011], or reducing the datapath size from 256-bits to 32-bits where our simulation results show $< 5%$ performance reduction).

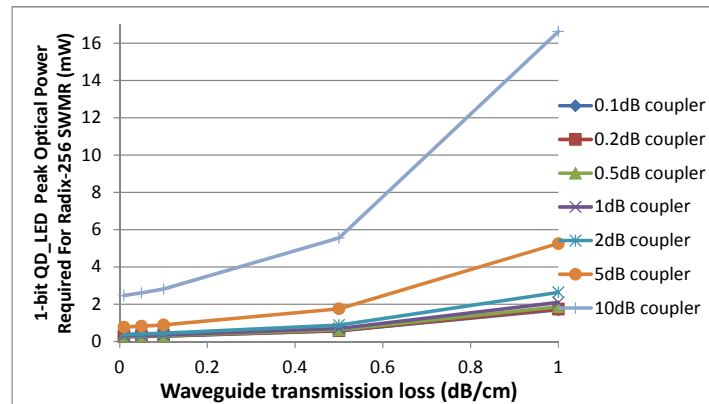
3.4. Qualitative Comparison to rNoC

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. Then we compare the process variation difference and the key device parameters of the two technologies.

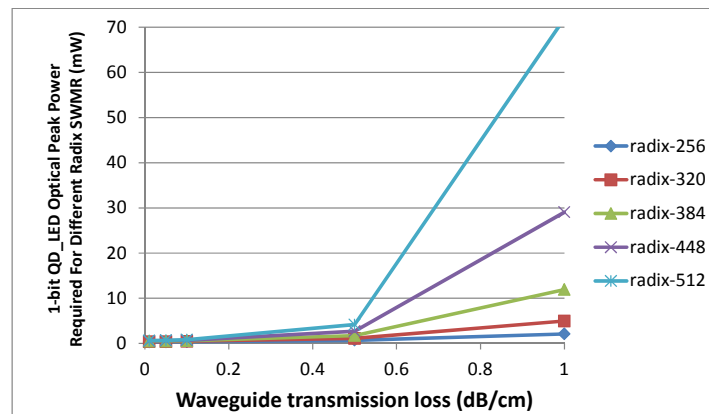
3.4.1. Scalability of mNoC. As mentioned in Section 2.2, rNoC scalability is limited by either nonlinear device behavior or ring trimming power. Both silicon waveguides and rings have nonlinear behavior with rings ($0.6mW$) generally entering nonlinearity at lower power values than silicon waveguides ($115mW$) [Biberman et al. 2011]. Silica waveguides in mNoC support much higher injected power (in Watts range) than silicon waveguides in rNoC. mNoCs do not utilize rings and neither QD LEDs or chromophores exhibit nonlinear behavior. Therefore, mNoC should be able to scale to higher radix.

The required output optical power from QD LEDs is a function of total waveguide transmission loss, optical devices' insertion loss (coupler, splitter etc. See more parameters in the Section 5) and the power required at each receiver (photodetector) on the worst case optical path (farthest from the source).

²The crosstalk is calculated in a way that for N evenly spaced emission spectrum of QD LED in $400-800nm$, go through the commercially available chromophore library to pick N chromophores with minimum crosstalk for neighboring channels



(a)



(b)

Fig. 4. 1-bit QD_LED Optical Power Required for (a) Radix-256 SWMR mNoC With Different Coupler Loss (b) SWMR mNoC for Different Radix with 1dB Coupler Loss

Figure 4(a) shows the peak optical power required for a 1-bit QD_LED in order to scale a SWMR mNoC crossbar to radix-256 with different waveguide transmission loss and coupler loss. The peak power does not consider the actual link utilization, and therefore for real workloads the required optical power is much lower. In the waveguide loss range 0.01dB/cm-1dB/cm, the QD_LED optical power stays far below $4mW$ (a reasonable output power a QD_LED supports) except for 5dB and 10dB coupler loss. We take 1dB coupler loss and calculate the required QD_LED peak optical power for even higher radix as shown in Figure 4(b). From this graph we see an opportunity to scale the crossbar to radix-512 especially when the transmission loss is below 0.5dB/cm.

3.4.2. Process Variation Comparison. One fundamental challenge for ring-based photonic technology is process variation which can cause ring resonators to malfunction. For QD_LEDs, although the size of all quantum dots for a desired wavelength might not be exactly the same, they still function correctly, but produce a broader full width half maximum (FWHM, 30-40nm). The average size of quantum dots affects the emission spectrum and it is an engineering process during fabrication to figure out the best method rather than a process variation problem which requires online tuning.

Gopal et al. [Gopal et al. 2009] successfully demonstrate a microcontact printing method for QD-LEDs to provide multi-color emission peak wavelengths. Different sizes of quantum dots can be stably achieved.

3.4.3. Key Parameter Comparison. 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 the farthest node, all the other inputs to nodes along the way should also be detected. The sensitivity varies from $80\mu W$ to $0.1\mu W$ [Pan et al. 2010; Li et al. 2011; Sahni et al. 2008] for photodetectors used on chip. To match the scalability analysis above, we choose a photodetector with $1\mu W$ sensitivity, which equals to about 30 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 300 photons are absorbed to achieve that sensitivity which requires 300 chromophores. If we assume one chromophore takes $2nm \times 2nm$ area, it is $1,200nm^2$ in total, which is negligible compared to the waveguide width. For the photodetector itself, the active area can be as small as $8\mu m^2$ [Feng et al. 2009].

Table I summarizes the key parameters for comparison of mNoC to ring resonator based networks (rNoC). The transmitter and receiver are within the same order of magnitude as a ring resonator, and the waveguide diameter is at least a quarter of the pitch size of that used in rNoC. The receiver area depends on the sensitivity of the photodetector, which determines the minimum input power for the farthest node from the source on the communication path.


Table I. Key Parameters of Ring Resonator NoC and mNoC

Component	Term	rNoC	mNoC
All	wavelength (nm)	1550	400-800
	trimming power(W)	(30)/1M rings	0
Transmitter	modulator type	ring modulator	QD-LED
	nonlinearity(mW)	0.6	none
	modulator size(μm^2)	7-100	200(256×256);19(64×64)
	modulation rate(GHz)	~ 10	1.1-67 (theoretical value)
	source efficiency(%)	30	~ 18
	source energy injection rate	1	0.5
Receiver	filter type	ring resonator	chromophore
	filter size(μm^2)	7-100	8
	sensitivity(μW)	1	1
Waveguide	#wavelengths	64	5-9
	size	$4\mu m$ (pitch)	$0.45-0.95\mu m$ (diameter)
	trans. loss(dB/cm)	0.2-1.5	0.05-1

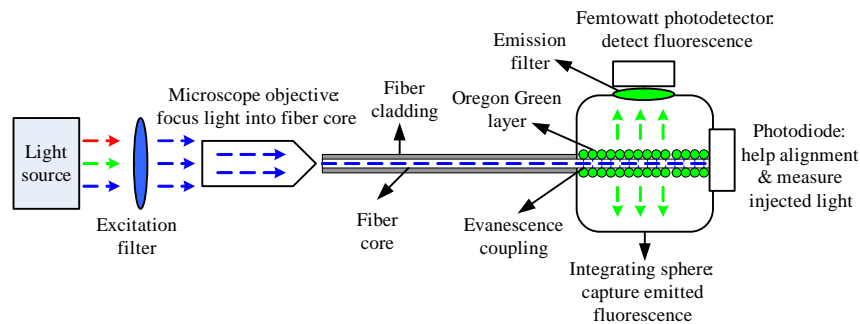
3.5. Fabrication and Experimental Demonstration

All the optical devices used in mNoC are compatible with current silicon fabrication technology [Liao et al. 2011; Sahni et al. 2007; Sahni et al. 2008; Arians et al. 2008; Wood and Bulović 2010; Kümmell et al. 2009; Kim et al. 2008]. 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 [Liao et al. 2011]. Therefore, we developed a prototype fiber-based system to demonstrate evanescent



(a)



(b)

Fig. 5. Experimental Setup (a) Real Setup (b) Schematic Structure

wave coupling of chromophores to a waveguide. Optical inputs with various wavelengths can be inserted into the same fiber by superposition (i.e., photons can occupy the same space) and each can selectively excite the desired chromophores if it is within 10-50nm of the fiber optic core by evanescent wave coupling. 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 $400\mu\text{m}$ core quartz fiber, and then use dip- or drop- coating to apply one or more layers of chromophores to the exposed core. To improve the coupling efficiency we also employ oxygen plasma to clean the stripped fiber ends. Upon injection of input photons the RET gates emit output photons radiating outward from the sidewall of the fiber core according to its transfer function. For non-cylinder silica waveguide, we can use similar coating technique to apply one or more layers of chromophores on the surface of waveguide. If precise control of the location of chromophores is required, we can use DNA self assembly fabrication method. The cost and overhead related to chromophore integration is very low [Dwyer and Lebeck 2007].

The experimental setup to inject input light and observe output is shown in Figure 5. 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




Fig. 6. Excitation Spectrum of Oregon GreenTM 488 from Gaussian Fit to the Measured Fluorescence Intensity and Vendor Data

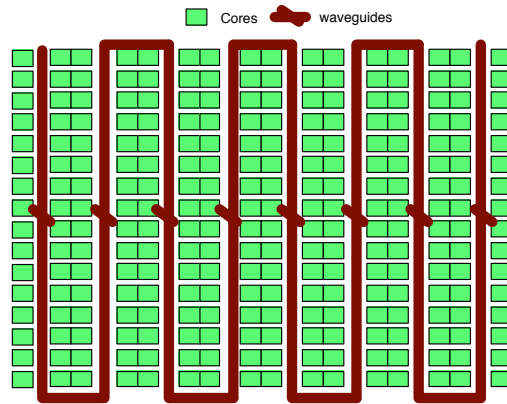


Fig. 7. Serpentine layout for 256×256 mNoC

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 6 shows two excitation spectrum curves of Oregon GreenTM 488. One is from Gaussian fit to the output fluorescence intensity excited by different wavelengths of light observed through a 532nm emission filter. The other is the expected result from vendor data. The two curves match each other well except a little peak shift due to Gaussian fit to data from our use of a limited number of excitation filters. This experiment demonstrates the successful evanescent coupling of chromophores to a waveguide.

4. MNOC ARCHITECTURE

Our proposed chromophore-based technology does not support switching functionality in the optical domain, therefore non-switch-based topologies, such as bus-based crossbars, may be more suitable. This section discusses the detailed design of the mNoC crossbar including the topology and packet design.

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. Our previous work discusses other topologies [Pang et al. 2013] (due to page limitations we omit these topologies). 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.

4.2. 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.

4.2.1. 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 trimming 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, the largest portion of the total area is from waveguides. For existing 64×64 bus-based crossbars, mNoC area is $1.44 \times$ compared to rNoC. However, this area easily fits into a $400mm^2$ chip using the 3D stack proposed for Corona [Vantrease et al. 2008].

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 width ($0.45\text{-}0.95\mu m$ in our case as shown in Table I); N_{nodes} is the number of nodes (256) connected by the waveguides; N_{col} is the number of columns of waveguides (8).

$$W_{total} = W_{pitch} \times \frac{datapath_size}{wavelength_density} \times N_{nodes} \times N_{col} \quad (1)$$

If the network datapath size is 256-bit, the total waveguide width is in the range of $4.7\text{-}10cm$ which is beyond the $2cm$ width of the $400mm^2$ die. However, there are three solutions to address this problem. First, use multi-layer of photonic fabrication and designs [Biberman et al. 2011]. 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. Third, reduce the datapath size of the network. With a 32-bit datapath, only a single layer integration is enough to fit the required waveguides.

Frequency: The switching frequency of the QD_LED has a theoretical range from 1.1 to $67GHz$, as shown in Table I, depending on the fluorescence lifetime of the selected

quantum dots. This frequency determines the fastest switching frequency, and thus data rate, in the network.

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 if a SWMR structure is used.

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 64×64 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 [Woo et al. 1995]. This includes performance and energy comparison, datapath impact and longevity discussion. Since Graphite [Miller et al. 2010] does not strictly enforce event ordering, our simulation data might not be precise. However, Graphite is the only non trace-based simulator we find which supports 256 or more cores.

5.1. Experimental Setup

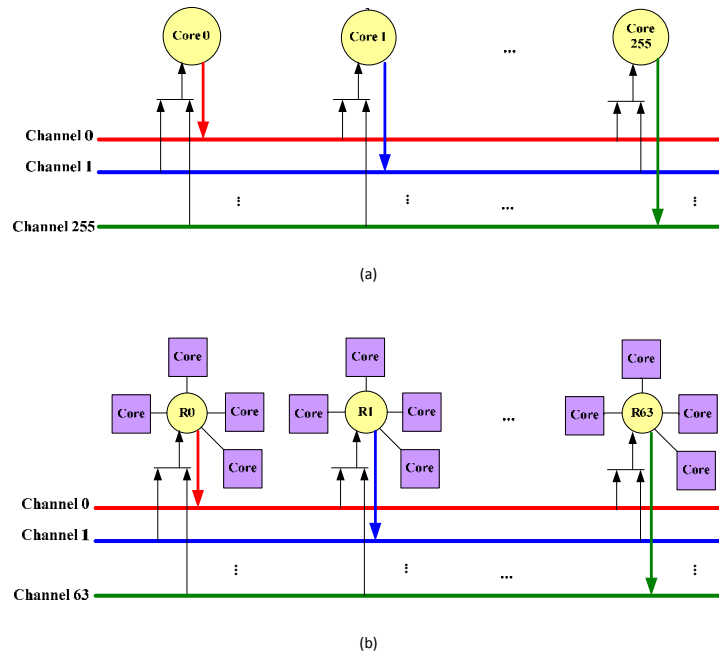


Fig. 8. Topology of Crossbars We Evaluate: (a)mNoC and (b) rNoC

The networks we evaluate are listed in Table II and each of them has 256 cores in total. The first is our radix-256 mNoC and the second is a radix-64 Firefly [Pan et al. 2009] - like network—rNoC as shown in Figure 8. We use a normal 4 stage

Table II. Networks Evaluated

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 III. Simulation configuration

Flit size	256-bit
Total buffer per link	1.5 KB (48 256-bit flit buffer)
Router pipeline stages	4 cycles
Electrical link latency	1 cycles
Optical link latency	1-9 cycles for mNoC; 1-5 for rNoC
Clock	5GHz
Core model	in-order model, private 32KB L1D, 32KB L1I, 512KB L2 Cache

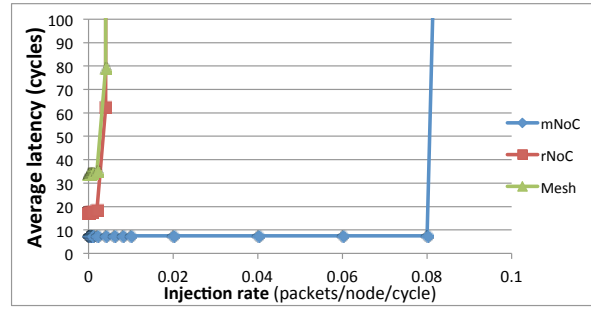
pipelined router for rNoC to connect its four cores in the cluster, and also to connect cores in the traditional electrical mesh network. We create all three networks in the Graphite [Miller et al. 2010] simulator and run all the simulations in the full simulator mode (vs lite mode with limited functionality [Miller et al. 2010]). The simulation configuration is summarized in Table III. 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 is approximately 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 9 cycles in the worst case. All the electrical links are modeled as 1 cycle [Chen et al. 2007] for the alternative networks.

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.

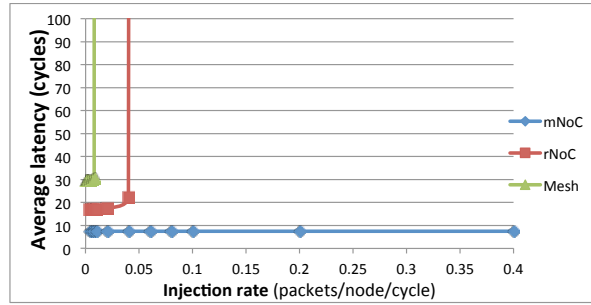
5.2. Evaluation with synthetic Benchmarks

We propose to use 1.5KB buffer/link (48 flit buffers with 256-bit per flit) which is the same as Firefly [Pan et al. 2009] and a simple flow-control mechanism which monitors the output buffers and links in mNoC. However, Graphite only supports infinite buffer without any flow-control mechanism. Therefore, we perform network only tests with three synthetic benchmarks: hotspot, uniform random, and tornado in GEM5 [Binkert et al. 2011a] with the same configuration to help better interpret Graphite’s results.

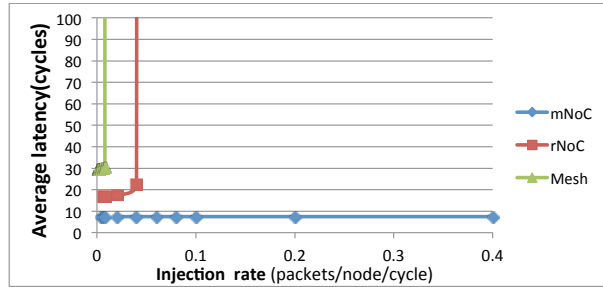
We use a network tester from GEM5 and the tester uses a dummy protocol in which all cache controllers send messages to directories while directories simply discard them. The average latency is measured from when a packet enters the outgoing queue of the source node until it is popped off the incoming queue at the destination node. mNoC has less than half the latency compared with rNoC in all three benchmarks as shown in Figure 9. Figure 9(a) shows average message latency versus injection rate (packets/node/cycle) of Hotspot traffic pattern for the three 256-node networks with a single directory. From Figure 9 we note that mNoC has lower average latency compared to the alternatives. Furthermore, when the injection rate increases, mNoC tolerates more Hotspot traffic. This is because the mNoC’s larger crossbar provides all to all connectivity and delivers messages efficiently. The overall injection rate of the three networks are low because the single directory is the bottleneck.



(a)



(b)



(c)

Fig. 9. Average Packet Latency for Synthetic traffic pattern (a)Hotspot, (b)Uniform Random, and (c)Tornado

We also run uniform random and tornado benchmarks with 16 directories and present similar results in Figure 9(b) and (c). Alternative networks saturate as injection rate increases, while mNoC sustains performance. Note that rNoC and Mesh saturate at a much higher rate here than in Hotspot because the directory number is $16\times$ bigger and thus the tolerance for traffic is also much higher. mNoC does not saturate because the buffers are big enough to accommodate all the packets. On average, mNoC has half the latency of rNoC and a quarter the latency of Mesh. In general, all the three networks perform better with more directories.

This GEM5 evaluation suggests that the 48 flit buffers provides the same performance as infinite buffers under light load. Our experiment results as presented later show that the link utilization is only 0.1% for 12 SPLASH benchmarks. Therefore, al-

though Graphite has infinite buffer size with no flow control, with low link utilization the performance evaluation data should still be valid.

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.

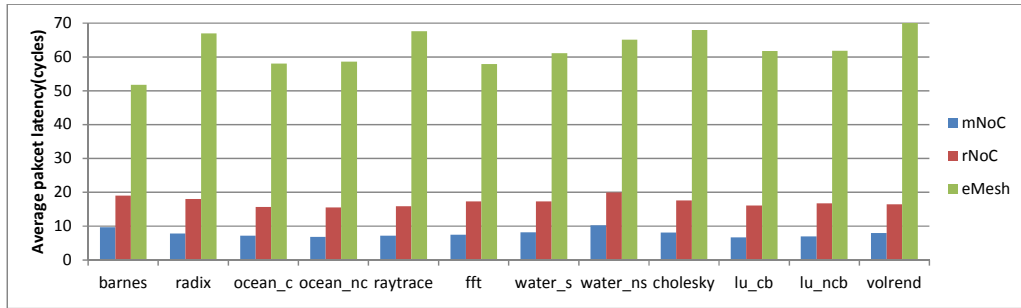


Fig. 10. Average Packet Latency of SPLASH Benchmarks

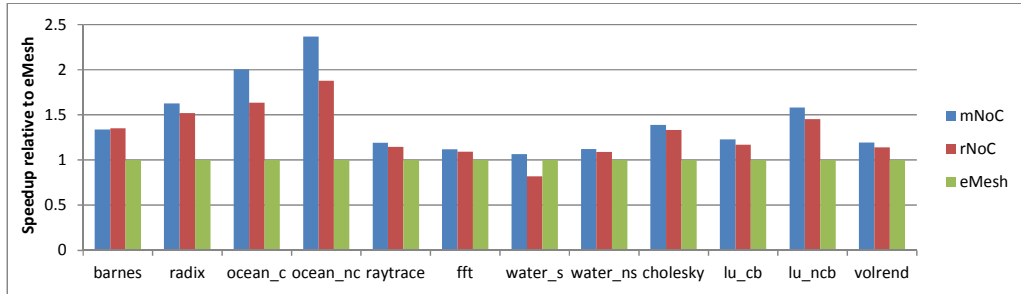


Fig. 11. Speedup Comparison of SPLASH Benchmarks

5.3.1. Performance Comparison. From the previous synthetic workloads, we see mNoC's performance advantage under certain traffic patterns: hotspot, uniform random, and tornado. 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[Woo et al. 1995], and run simulations with 256 threads in Graphite. Figure 10 shows the average packet latency for each of the three networks and Figure 11 shows speedup relative to mesh.

The mNoC crossbar shows significant average latency advantage over the other two alternatives in Figure 10. 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

Table IV. Optical energy parameters

QD_LED energy efficiency	18%
QD_LED 1-to-0 ratio	1
Waveguide loss	1dB/cm
Coupler loss	1dB
Power loss of chromophores	0.5 μ W for 1 μ W sensitivity
Optical splitter	0.2dB
photodetector	0.1dB
Link utilization(SPLASH)	0.1%

a flattened structure and it can efficiently transmit packets without taking multiple hops. The improvement of packet latency also leads to better overall system performance. Figure 11 shows 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 [Barrow-Williams et al. 2009], 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 IV. To match the clock rate, we evaluate the quantum dots with lifetime of 200ps [Kümmell et al. 2009; Ariens et al. 2008] in our model which corresponds to signal switching frequency of 5GHz in the network. We assume the source energy efficiency is 18%³.




Fig. 12. Link utilization of SPLASH Benchmarks

A notable feature of mNoC is that the QD_LED light source 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. From running all the 12 SPLASH benchmarks, we gathered the link utilization of mNoC as shown in Figure 12. On average the link utilization is 0.1%, which is low due to limited number of shared memory accesses [Barrow-Williams et al. 2009], abundant on-chip resources (e.g., caches), the application characteristics and the high speed of the optical communication structure.

³From QD Vision, <http://www.qdvision.com/content1599>

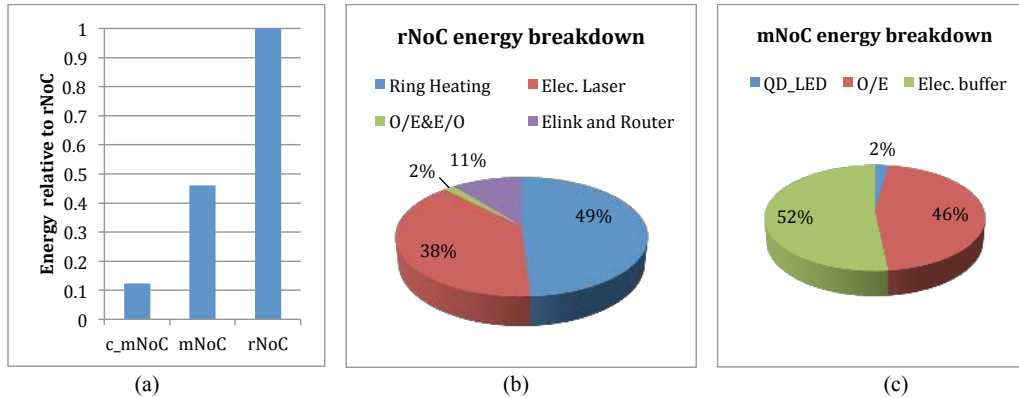


Fig. 13. Energy comparison (a) average energy consumption across all SPASH benchmarks, (b) radix-256 mNoC, and (c) radix-64 clustered rNoC average energy breakdown.

We conservatively make the following assumptions for ring trimming power calculation to favor rNoC technology: 1) the total trimming power over 20K temperature range for 1 million rings in rNoC is the same as the average trimming power (49.4W) for 524K rings in Nitta, et al [Nitta et al. 2011], and 2) a process variation tolerance technique [Xu et al. 2012] is used to further reduce the trimming power by 39%, which is 30W in total for all the rings. We also ignore the scattering loss from off-resonance rings and optimistically assume that there are splitters before ring receivers to extract the exact amount of power to minimize insertion loss caused by serial ring resonators. Chromophores only consume power when excited and the power loss is $0.5\mu W$ for $1\mu W$ sensitivity photodetector. We use 5 wavelengths per waveguide in the evaluation and the crosstalk is -10dB for the two neighbouring channels. The remaining optical parameters for ring-based devices are the same as Flexishare [Pan et al. 2010].

For electrical circuit and links, we estimate the static power for buffers in mNoC using Cacti 6.5 ITRS-HP conservative 32nm technology, which is the main electrical power consumption in mNoC. The rest of the electrical parameters, we refer to models presented elsewhere [Pan et al. 2010; Pan et al. 2009; Joshi et al. 2009]. We add one more network—clustered_mNoC(c_mNoC) for comparison. c_mNoC has the same organization as rNoC, but instead of using ring resonator-based technology, we use the proposed molecular scale technology.

Figure 13 presents the average energy results across all 12 SPLASH benchmarks. The source energy is directly proportional to link utilization as shown in Figure 12 and individual benchmarks will incur differing energy accordingly. For brevity we include only the average here. Figure 13(a) shows that the average energy of a radix-256 mNoC crossbar is reduced by 54% (the worst case for benchmark radix with the highest link utilization is 50%) compared to the radix-64 clustered rNoC (256 total cores). This improvement is due to removal of ring resonator trimming power as shown in the energy breakdown graph Figure 13(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 13(a) uses only 12% of the energy as rNoC, which has a similar structure but uses different technology.

5.3.3. Datapath Width Impact. A wide datapath is efficient for utilizing on-chip bandwidth, however, it also increases cost such as area as analyzed in Section 4.2.1. We

compare the performance and energy of SPLASH benchmarks with different NoC datapath widths and the results are shown in Figure 14. Figure 14(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 serialization latency of a smaller datapath. Figure 14(b) shows the smaller datapath width can efficiently reduce energy by up to 85% for width 32.




Fig. 14. Performance (a) and average energy across all SPLASH benchmarks (b) comparison of different datapath widths

5.3.4. Longevity. Despite the small size and low power consumption of chromophores, they also have a photobleaching problem [Lakowicz 2006] 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 [Ow et al. 2005; Canton et al. 2011].

A single chromophore takes $4nm^2$ area and we allocate as many chromophores as possible into the $0.625\mu m^2$ area 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 [Langhals 2005; Langhals 1995; Langhals et al. 1998]. If we consider the chromophores utilization and also use the encapsulation to extend the longevity by a factor of 10 (F_{encap}) [Canton

et al. 2011], 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.

$$N_{years} = \frac{N_{chromophores} \times L_{single}}{N_{year_cycles} \times U_{link} \times P_{1_to_0}} \times F_{encap} \quad (2)$$

6. ARCHITECTURAL IMPLICATIONS

With respect to power, molecular optical technology enables us to reduce the on-chip energy 87% by keeping the same radix of the crossbar, or scale to 256×256 SWMR crossbar with 54% 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 [Kim et al. 2005]. 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 [Binkert et al. 2011b].

Molecular technology introduced by this paper is a good candidate to build a high radix router. The 256×256 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 4.

6.2. Cache Coherence

One feature of our SWMR 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 future work.

7. CONCLUSION

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 ring trimming, 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 on-chip emitters such as 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. Furthermore, mNoC trades static energy for dynamic energy and greatly reduces energy consumption. An mNoC can achieve 88% reduction in energy for a 64×64 crossbar compared to similar ring resonator based designs. Additionally, an mNoC can scale to a 256×256 crossbar with 54% 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. Fourth, explore additional protocols to reduce receiver buffering requirements.

References

- PO Anikeeva, CF Madigan, JE Halpert, MG Bawendi, and V. Bulović. 2008. Electronic and excitonic processes in light-emitting devices based on organic materials and colloidal quantum dots. *Physical Review B* 78, 8 (2008), 085434.
- R. Arians, A. Gust, T. Kummell, C. Kruse, S. Zaitsev, G. Bacher, and D. Hommel. 2008. Electrically driven single quantum dot emitter operating at room temperature. *Applied Physics Letters* 93, 17 (2008), 173506–173506.
- N. Barrow-Williams, C. Fensch, and S. Moore. 2009. A communication characterisation of Splash-2 and Parsec. In *Workload Characterization, 2009. IISWC 2009. IEEE International Symposium on*. IEEE, 86–97.
- A. Biberman, K. Preston, G. Hendry, N. Sherwood-Droz, J. Chan, J.S. Levy, M. Lipson, and K. Bergman. 2011. Photonic network-on-chip architectures using multilayer deposited silicon materials for high-performance chip multiprocessors. *ACM Journal on Emerging Technologies in Computing Systems (JETC)* 7, 2 (2011), 7.
- N. Binkert, B. Beckmann, G. Black, S.K. Reinhardt, A. Saidi, A. Basu, J. Hestness, D.R. Hower, T. Krishna, S. Sardashti, and others. 2011a. The gem5 simulator. *ACM SIGARCH Computer Architecture News* 39, 2 (2011), 1–7.
- N. Binkert, A. Davis, N.P. Jouppi, M. McLaren, N. Muralimanohar, R. Schreiber, and J.H. Ahn. 2011b. The role of optics in future high radix switch design. In *Proceeding of the 38th annual international symposium on Computer architecture*. ACM, 437–448.
- G. Canton, R. Ricco, F. Marinello, S. Carmignato, and F. Enrichi. 2011. Modified Stöber synthesis of highly luminescent dye-doped silica nanoparticles. *Journal of Nanoparticle Research* 13, 9 (2011), 4349–4356.
- J. Chan, G. Hendry, K. Bergman, and L. P. Carloni. 2011. Physical-Layer Modeling and System-Level Design of Chip-Scale Photonic Interconnection Networks. *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on* 30, 10 (2011), 1507–1520.
- G. Chen, H. Chen, M. Haurylau, N.A. Nelson, D.H. Albonese, P.M. Fauchet, and E.G. Friedman. 2007. Predictions of CMOS compatible on-chip optical interconnect. *Integration, the VLSI journal* 40, 4 (2007), 434–446.

- S. Chen, L. Zhang, Y. Fei, and T. Cao. 2012. Bistability and self-pulsation phenomena in silicon microring resonators based on nonlinear optical effects. *Optics Express* 20, 7 (2012), 7454–7468.
- M. J. Cianchetti, J. C. Kerekes, and D. H. Albonesi. 2009. Phastlane: a rapid transit optical routing network. In *36th annual international symposium on Computer architecture*. 441–450.
- C. Dwyer and A. R. Lebeck. 2007. *Introduction to DNA Self-Assembled Computer Design*. Artech House, Inc., Norwood, MA, USA.
- D. Feng, S. Liao, P. Dong, N. N. Feng, H. Liang, D. Zheng, C. C. Kung, J. Fong, R. Shafiiha, J. Cunningham, and others. 2009. High-speed Ge photodetector monolithically integrated with large cross-section silicon-on-insulator waveguide. *Applied Physics Letters* 95 (2009), 261105.
- A. Gopal, K. Hoshino, S. Kim, and X. Zhang. 2009. Multi-color colloidal quantum dot based light emitting diodes micropatterned on silicon hole transporting layers. *Nanotechnology* 20, 23 (2009), 235201.
- F. Hargart, C. A. Kessler, T. Schwarzback, E. Koroknay, S. Weidenfeld, M. Jetter, and P. Michler. 2013. Electrically driven quantum dot single-photon source at 2 GHz excitation repetition rate with ultra-low emission time jitter. *Applied Physics Letters* 102, 1 (2013), 011126–011126.
- A. Joshi, C. Batten, Y.J. Kwon, S. Beamer, I. Shamim, K. Asanovic, and V. Stojanovic. 2009. Silicon-photonics networks for global on-chip communication. In *Proceedings of the 2009 3rd ACM/IEEE International Symposium on Networks-on-Chip*. IEEE Computer Society, 124–133.
- J. Kim, W.J. Dally, B. Towles, and A.K. Gupta. 2005. Microarchitecture of a high-radix router. In *ACM SIGARCH Computer Architecture News*, Vol. 33. IEEE Computer Society, 420–431.
- K. H. Kim, D. L. Nguyen, H. Lim, P. T. Nga, Y. H. Cho, and others. 2011. Shell layer dependence of photoblinking in CdSe/ZnSe/ZnS quantum dots. *Applied Physics Letters* 98 (2011), 012109.
- L. A. Kim, P. O. Anikeeva, S. A. Coe-Sullivan, J. S. Steckel, M. G. Bawendi, and V. Bulovic. 2008. Contact printing of quantum dot light-emitting devices. *Nano letters* 8, 12 (2008), 4513–4517.
- N. Kirman, M. Kirman, R.K. Dokania, J.F. Martinez, A.B. Apse, M.A. Watkins, and D.H. Albonesi. 2006. Leveraging optical technology in future bus-based chip multiprocessors. In *Proceedings of the 39th Annual IEEE/ACM International Symposium on Microarchitecture*. IEEE Computer Society, 492–503.
- T. Kümmell, R. Ariens, A. Gust, C. Kruse, S. Zaitsev, D. Hommel, and G. Bacher. 2009. Electrically driven room temperature operation of a single quantum dot emitter. In *Proceedings of SPIE*, Vol. 7211. 72110G.
- J.R. Lakowicz. 2006. *Principles of fluorescence spectroscopy*. Vol. 1. Springer.
- H. Langhals. 1995. Cyclic Carboxylic Imide Structures as Structure Elements of High Stability Novel Developments in Perylene Dye Chemistry. *Heterocycles-Sendai Institute of Heterocyclic Chemistry* 40, 1 (1995), 477.
- H. Langhals. 2005. Control of the Interactions in Multichromophores: Novel Concepts. Perylene Bis-imides as Components for Larger Functional Units. *Helvetica chimica acta* 88, 6 (2005), 1309–1343.
- H. Langhals, J. Karolin, and L.B.Å. Johansson. 1998. Spectroscopic properties of new and convenient standards for measuring fluorescence quantum yields. *Journal of the Chemical Society, Faraday Transactions* 94, 19 (1998), 2919–2922.
- S. Le Beux, I. O'Connor, G. Nicolescu, G. Bois, and P. Paulin. 2013. Reduction methods for adapting optical network on chip topologies to 3D architectures. *Microprocessors and Microsystems* 37, 1 (2013), 87–98.
- S. Le Beux, J. Trajkovic, I. O'Connor, G. Nicolescu, G. Bois, and P. Paulin. 2010. Multi-Optical Network-on-Chip for Large Scale MPSoC. *Embedded Systems Letters, IEEE* 2, 3 (2010), 77–80.
- Z. Li, D. Fay, A. Mickelson, L. Shang, M. Vachharajani, D. Filipovic, W. Park, and Y. Sun. 2009. Spectrum: A hybrid nanophotonic electric on-chip network. In *Design Automation Conference, 2009. DAC'09. 46th ACM/IEEE*. IEEE, 575–580.
- Z. Li, M. Mohamed, X. Chen, H. Zhou, A. Mickelson, L. Shang, and M. Vachharajani. 2011. Iris: A hybrid nanophotonic network design for high-performance and low-power on-chip communication. *ACM Journal on Emerging Technologies in Computing Systems (JETC)* 7, 2 (2011), 8.
- S. Liao, N.N. Feng, D. Feng, P. Dong, R. Shafiiha, C. C. Kung, H. Liang, W. Qian, Y. Liu, J. Fong, J. E. Cunningham, Y. Luo, and M. Asghari. 2011. 36 GHz submicron silicon waveguide germanium photodetector. *Optics Express* 19, 11 (2011), 10967–10972.
- O. Liboiron-Ladouceur, I. Cerutti, P.G. Raponi, N. Andriolli, and P. Castoldi. 2011. Energy-efficient design of a scalable optical multiplane interconnection architecture. *Selected Topics in Quantum Electronics, IEEE Journal of* 17, 2 (2011), 377–383.
- O. Liboiron-Ladouceur, A. Shacham, B. A. Small, B. G. Lee, H. Wang, C. P. Lai, A. Biberman, and K. Bergman. 2008. The data vortex optical packet switched interconnection network. *Journal of Lightwave Technology* 26, 13 (2008), 1777–1789.

- B. S. Mashford, M. Stevenson, Z. Popovic, C. Hamilton, Z. Zhou, C. Breen, J. Steckel, V. Bulovic, M. Bawendi, S. Coe-Sullivan, and others. 2013. High-efficiency quantum-dot light-emitting devices with enhanced charge injection. *Nature Photonics* (2013).
- J. E. Miller, H. Kasture, G. Kurian, C. Gruenwald, N. Beckmann, C. Celio, J. Eastep, and A. Agarwal. 2010. Graphite: A distributed parallel simulator for multicores. In *High Performance Computer Architecture (HPCA), 2010 IEEE 16th International Symposium on*. IEEE, 1–12.
- C. Nitta, M. Farrens, and V. Akella. 2011. Addressing system-level trimming issues in on-chip nanophotonic networks. In *High Performance Computer Architecture (HPCA), 2011 IEEE 17th International Symposium on*. IEEE, 122–131.
- H. Ow, D.R. Larson, M. Srivastava, B.A. Baird, W.W. Webb, and U. Wiesner. 2005. Bright and stable core-shell fluorescent silica nanoparticles. *Nano letters* 5, 1 (2005), 113–117.
- Y. Pan, J. Kim, and G. Memik. 2010. Flexishare: Channel sharing for an energy-efficient nanophotonic crossbar. In *High Performance Computer Architecture (HPCA), 2010 IEEE 16th International Symposium on*. IEEE, 1–12.
- Y. Pan, P. Kumar, J. Kim, G. Memik, Y. Zhang, and A. Choudhary. 2009. Firefly: Illuminating future network-on-chip with nanophotonics. In *In Proc. of the International Symposium on Computer Architecture*.
- J. Pang, C. Dwyer, and A. R. Lebeck. 2013. Exploiting emerging technologies for nanoscale photonic networks-on-chip. In *Proceedings of the Sixth International Workshop on Network on Chip Architectures*. ACM, 53–58.
- H. Park, Y. Kuo, A.W. Fang, R. Jones, O. Cohen, M.J. Paniccia, and J.E. Bowers. 2007a. A hybrid AlGaInAs-silicon evanescent preamplifier and photodetector. *Opt. Express* 15, 21 (2007), 230–232.
- I. K. Park, M. K. Kwon, J. O. Kim, S. B. Seo, J. Y. Kim, J. H. Lim, S. J. Park, and Y. S. Kim. 2007b. Green light-emitting diodes with self-assembled In-rich InGaN quantum dots. *Applied Physics Letters* 91 (2007), 133105.
- G. Priem, P. Dumon, W. Bogaerts, D. Van Thourhout, G. Morthier, and R. Baets. 2006. Nonlinear effects in ultrasmall silicon-on-insulator ring resonators. 6183 (apr 2006).
- L. Ramini, P. Grani, H.T. Fankem, A. Ghiribaldi, S. Bartolini, and D. Bertozzi. 2014. Assessing the energy break-even point between an optical NoC architecture and an aggressive electronic baseline. In *Proceedings of the conference on Design, Automation & Test in Europe*. European Design and Automation Association, 308.
- K. Rurack and M. Spieles. 2011. Fluorescence Quantum Yields of a Series of Red and Near-Infrared Dyes Emitting at 600- 1000 nm. *Analytical chemistry* 83, 4 (2011), 1232–1242.
- S. Sahni, X. Luo, J. Liu, Y. Xie, and E. Yablonovitch. 2008. Junction field-effect-transistor-based germanium photodetector on silicon-on-insulator. *Opt. Lett.* 33 (May 2008), 1138–1140.
- S. Sahni, E. Yablonovitch, J. Liu, and Y. Xie. 2007. Germanium-on-SOI photo-detector based on an FET structure. In *CLEO*. Optical Society of America.
- L. Tong, R. R. Gattass, J. B. Ashcom, S. He, J. Lou, M. Shen, I. Maxwell, and E. Mazur. 2003. Subwavelength-diameter silica wires for low-loss optical wave guiding. *Nature* 426, 6968 (2003), 816–819.
- B. Valeur. 2001. Molecular fluorescence: principles and applications. (2001).
- D. Vantrease, R. Schreiber, M. Monchiero, M. McLaren, N. P. Jouppi, M. Fiorentino, A. Davis, N. Binkert, R. G. Beausoleil, and J. H. Ahn. 2008. Corona: System implications of emerging nanophotonic technology. In *Proceedings of the 35th Annual International Symposium on Computer Architecture*. IEEE Computer Society, Washington, DC, USA, 153–164.
- J. Wang, W. Y. Loh, K. T. Chua, H. Zang, Y. Z. Xiong, T. H. Loh, M. B. Yu, S. J. Lee, G. Q. Lo, and D. L. Kwong. 2008. Evanescent-coupled Ge pin photodetectors on Si-waveguide with SEG-Ge and comparative study of lateral and vertical pin configurations. *Electron Device Letters, IEEE* 29, 5 (2008), 445–448.
- S. C. Woo, M. Ohara, E. Torrie, J. P. Singh, and A. Gupta. 1995. The SPLASH-2 programs: characterization and methodological considerations. In *Proceedings of the 22nd annual international symposium on Computer architecture*. ACM, 24–36.
- V. Wood and V. Bulović. 2010. Colloidal quantum dot light-emitting devices. *Nano Reviews* 1, 0 (2010).
- Y. Xu, J. Yang, and R. Melhem. 2012. Tolerating process variations in nanophotonic on-chip networks. In *Proceedings of the 39th International Symposium on Computer Architecture*. IEEE Press, 142–152.
- Y. Ye, J. Xu, X. Wu, W. Zhang, W. Liu, and M. Nikdast. 2012. A Torus-Based Hierarchical Optical-Electronic Network-on-Chip for Multiprocessor System-on-Chip. *ACM Journal on Emerging Technologies in Computing Systems (JETC)* 8, 1 (2012), 5.
- P. P. Yupapin, C. Teeka, and P. Chitsakul. 2006. Mathematical Simulation of Nonlinear Effects in Micro Ring Resonator. In *Emerging Technologies-Nanoelectronics, 2006 IEEE Conference on*. IEEE, 316–321.