

# More is Less, Less is More: Molecular-Scale Photonic NoC Power Topologies

Jun Pang

Department of Computer Science  
Duke University  
pangjun92@gmail.com

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

Molecular-scale Network-on-Chip (mNoC) crossbars use quantum dot LEDs as an on-chip light source, and chromophores to provide optical signal filtering for receivers. An mNoC reduces power consumption or enables scaling to larger crossbars for a reduced energy budget compared to current nanophotonic NoC crossbars. Since communication latency is reduced by using a high-radix crossbar, minimizing power consumption becomes a primary design target. Conventional Single Writer Multiple Reader (SWMR) photonic crossbar designs broadcast all packets, and incur the commensurate required power, even if only two nodes are communicating.

This paper introduces power topologies, enabled by unique capabilities of mNoC technology, to reduce overall interconnect power consumption. A power topology corresponds to the logical connectivity provided by a given power mode. Broadcast is one power mode and it consumes the maximum power. Additional power modes consume less power but allow a source to communicate with only a statically defined, potentially non-contiguous, subset of nodes. Overall interconnect power is reduced if the more frequently communicating nodes use modes that consume less power, while less frequently communicating nodes use modes that consume more power. We also investigate thread mapping techniques to fully exploit power topologies. We explore various mNoC power topologies with one, two and four power modes for a radix-256 SWMR mNoC crossbar. Our results show that the combination of power topologies and intelligent thread mapping can reduce total mNoC power by up

to 51% on average for a set of 12 SPLASH benchmarks. Furthermore performance is 10% better than conventional resonator-based photonic NoCs and energy is reduced by 72%.

**Categories and Subject Descriptors** C.1.2 [Multiple Data Stream Architectures (Multiprocessors)]: Interconnection architectures

**Keywords** Nanophotonics; Interconnection Network; Energy Efficiency

## 1. Introduction

Nanophotonic Network-on-Chips (NoC) exhibit superior power delay product and bandwidth compared to electrical NoCs [19, 27, 28, 36] and are likely to play an important role in next generation NoCs. The ideal NoC is one large crossbar that enables communication between all pairs of nodes in the absence of output conflicts. Unfortunately, current ring resonator based crossbars are difficult to scale larger than  $64 \times 64$  due to ring thermal tuning, ring nonlinearity, and external light source inefficiency.

Recently proposed Molecular-scale Network-on-Chip (mNoC) designs use emerging molecular-scale devices to construct nanophotonic networks [29, 30]. Quantum dot LEDs (QD\_LED) provide on-chip electrical to optical signal generation/modulation and chromophores provide optical signal filtering for receivers. These molecular devices replace the ring resonators and external laser source used in current nanophotonic NoCs. For a simple Single Writer Multiple Reader (SWMR) crossbar, mNoC reduces energy consumption or enables scaling to larger crossbars (more than radix-256) for a smaller energy budget and with better performance compared to ring resonator NoCs.

Communication latency and power consumption are two important factors in NoC design. Communication latency is reduced in high radix (e.g.,  $256 \times 256$ ) SWMR mNoC crossbars since there are no intermediate routers and optical signals transmit at the speed of light. Therefore, power consumption becomes a top concern in mNoC design. In con-

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

ASPLOS '15, March 14–18, 2015, Istanbul, Turkey.  
Copyright © 2015 ACM 978-1-4503-2835-7/15/03...\$15.00.  
<http://dx.doi.org/10.1145/2694344.2694377>

ventional SWMR crossbars, a source node broadcasts every packet and consumes power necessary to reach every destination even with only one or a few destination nodes. Furthermore, broadcast ignores frequency of communication between a source and specific destination nodes. The challenge is to design a single physical interconnect that provides the latency benefits of high-radix crossbars, but without the worst case broadcast power consumption for frequent communication to subsets of destinations.

To meet the above challenge, this paper introduces mNoC power topologies, a unique capability where connectivity is dynamically controlled as a function of source power (**Contribution 1**). For a given source on a single waveguide, broadcast requires the highest power since it reaches all destinations, whereas lower power modes provide connectivity to progressively smaller subsets of destination nodes. The ability to control the on-chip QD\_LED light source power and the appropriate design of optical waveguides enables different power modes. In particular, we use asymmetric waveguide splitters to allow potentially non-contiguous destinations in each power mode.


This paper explores static power topologies with a fixed set of source power modes. The specific power mode is selected dynamically based on destination. Our designs are guided by examining program communication patterns. The general design problem is NP-hard, therefore we use heuristic methods to design several power topologies with one, two, and four power modes (**Contribution 2**). We also explore the impact of thread mapping, since overall power is reduced if frequently communicating nodes use the lowest power modes, while less frequently communicating nodes use higher power modes (**Contribution 3**).

We use Graphite [25] to evaluate performance and to gather execution traces of several SPLASH-2 benchmarks [37] executing on a 256 core system. The traces are used to examine communication patterns and to compute NoC power consumption. Our results show that total mNoC power consumption is reduced by around 13% for naive power topologies, whereas a communication-aware power topology that utilizes intelligent thread mapping reduces power by 51%. The best overall design achieves performance 10% better than conventional resonator-based photonic NoCs while reducing energy by 72% (**Contribution 4**).

The remainder of this paper is organized as follows. Section 2 introduces background information for mNoC and the motivation of our work. We present power topologies and their implementation in Section 3 and Section 4 discusses architecting power topologies. In Section 5 we evaluate several power topologies and the impact of thread mapping. Related work is discussed in Section 6 and we conclude in Section 7.

## 2. Background and Motivation

Nanophotonic NoCs are promising since they provide high bandwidth and lower power than conventional electrical



**Figure 1.** mNoC Main Components

NoCs. Previously proposed ring resonator-based NoCs (rNoC) have the following main components: 1) an external laser source to provide highly coherent light, 2) a waveguide to provide connectivity, and 3) ring resonators for both modulation and detection. Despite the advantages over electrical NoCs, there still exist significant limitations to resonator-based NoCs, including: 1) high static power consumption due to the power inefficiency from the activity independent off-chip laser source and the non-negligible ring thermal tuning power ( $20\text{-}100\mu\text{W}/\text{ring}$  over 20K temperature range) [19, 26, 27], and 2) poor scalability due to ring resonators' nonlinearity [5]. Recently proposed molecular-scale nanophotonic network on-chip (mNoC) seeks to overcome these limitations [29, 30].

### 2.1 Enabling Molecular-Scale Nanophotonic Technology

Our work on power topologies builds on the base mNoC designs, but is applicable to any similar technology (e.g., VCSELs as light sources). We provide a brief summary of mNoC technology and a comparison to rNoC, further details are available elsewhere [29, 30]. Figure 1 shows the main mNoC components, which includes QD\_LEDs, chromophores and waveguides. Transmitters are composed of silicon compatible QD\_LEDs [13] 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, and have small size, narrow emission bandwidth, good stability and fast excitation rate [14, 24, 33, 38]. Receivers are composed from chromophores [23, 35] that filter the desired optical signals and couple the energy from the waveguide to a photodetector for O/E conversion. Subwavelength-diameter silica ( $\text{SiO}_2$ ) waveguides [34] transmit visible light from the transmitters to receivers. These components have been individually experimentally demonstrated [1, 13, 14, 20, 30, 34, 38] and are silicon compatible, however a fully integrated mNoC has not yet been demonstrated.

Table 1 summarizes the key differences between mNoC and rNoC for a 256-node system. In this table mNoC uses a single writer multiple reader (SWMR)  $256 \times 256$  crossbar with a serpentine layout whereas rNoC uses a clustered topology with a  $64 \times 64$  optical crossbar and 4 nodes per crossbar port (a clustered mNoC is also possible). Communication within a 4-node cluster uses the electrical domain.

Metric of Interest	rNoC	mNoC
<b>Technology Characteristics</b>		
Wavelength (nm)	1550	390-750
Requires Thermal Tuning	Yes	No
Activity/traffic independent light source	Yes	No
Nonlinearity (transmitters & receivers)	Yes	No
<b>System Evaluation Metrics</b>		
Scalability (Max crossbar size)	64×64	> 256×256
Normalized energy (256-node)	1	< 0.51
Normalized performance (256-node)	1	1.1


**Table 1.** Comparison between rNoC and mNoC

- **Scalability:** rNoC scalability is limited by either nonlinear device behavior or ring thermal tuning power. Since mNoCs do not utilize rings and neither QD\_LEDs [31] or chromophores exhibit nonlinear behavior, waveguide nonlinearity is the primary limitation to mNoC scalability. An mNoC crossbar can easily scale to more than radix-256 even with a  $2dB/cm$  loss waveguide. With multilayer silicon integration [5] even higher scalability may be achievable.
- **Energy and Performance:** Overall energy consumption is reduced because mNoC eliminates the large ring thermal tuning power. Moreover, the off-chip activity-independent laser source is replaced with on-chip QD\_LEDs, where the waveguide link utilization and per packet 1-to-0 ratio play an important role in further reducing energy consumption. Previous work [29, 30] shows that for a 256-node clustered design with a radix-64 optical crossbar, the clustered mNoC (c\_mNoC) reduces energy by 77% for the same performance compared to rNoC. When scaled to a single  $256 \times 256$  crossbar (not possible for rNoC), mNoC achieves a 49% energy reduction with a 10% increase in performance vs. the clustered rNoC topology.

## 2.2 Motivation

The above results demonstrate the potential benefit of mNoC technology over rNoC technology by providing more scalability (and thus performance) with lower energy consumption. We also observe that, for a 256-node system, c\_mNoC achieves the largest energy reduction while the full  $256 \times 256$  mNoC crossbar achieves the highest performance. Although it is possible to construct multiple networks to provide either low power or high performance when needed (e.g., catnap [11]), *ideally we want a single physical network that provides the performance of the high radix crossbar (mNoC) with the energy savings of the lower-radix clustered topology (c\_mNoC).*

Our work toward achieving this goal is motivated by three primary observations: 1) source power consumption dominates overall mNoC power consumption, 2) waveguide



**Figure 2.** Percent of mNoC Power for QD\_LED and O/E

loss determines mNoC source power requirements, and 3) workloads exhibit non-uniform communication patterns.


**Observation 1:** Source power dominates overall mNoC power consumption. The total power consumption of an mNoC is composed of three parts: QD\_LED source power, O/E conversion power due to photoreceiver power dissipation, and electrical circuit power (buffers, etc.). The O/E power is a function of the photoreceiver’s minimum input optical power (mIOP) [8, 20]; a low mIOP requires a high gain photoreceiver and thus increases O/E power while high mIOP can reduce O/E power but requires higher source power [5, 6]. This relationship enables mNoC designs to tradeoff O/E power for source power while keeping total power relatively constant. The impact of this tradeoff for rNoC is minimal since thermal tuning dominates total power.

Figure 2 shows how the QD\_LED and O/E power changes as a percentage of the total power as mIOP increases from  $1\mu W$  to  $10\mu W$ , assuming O/E conversion power decreases linearly with mIOP<sup>1</sup>. From this figure, we see that as mIOP increases, the dominant power is shifted from O/E to QD\_LED source power. Using a  $10\mu W$  mIOP photoreceiver, QD\_LED source power is 80% of the total power and becomes the new optimization target.

**Observation 2:** Waveguide loss determines source power required to reach a particular destination for a given photodetector mIOP. The waveguide loss dominates other losses due to chromophore receivers which are dependent on the specific mIOP for a photodetector. The two main sources of waveguide loss for mNoC are transmission loss given the distance to a destination and various insertion loss along the optical path (e.g., splitter and photodetector insertion loss). Consider the distance dependence loss; Figure 3 shows that mNoC source power increases exponentially vs. maximum broadcast distance to a destination on the waveguide.

**Observation 3:** Applications exhibit non-uniform communication patterns. Analysis of the SPLASH benchmarks shows that not every packet needs to traverse the full broadcast range. The average communication distance between

<sup>1</sup> Linearity is achieved by changing the number of gain stages and/or tuning the photoreceiver’s bias voltage ( $V_{bias}$ ).



**Figure 3.** Source Power Consumption vs. Broadcast Distance

threads, based on thread ID numbered from 0 to 255, is 102 across 12 SPLASH benchmarks. Therefore, broadcasting to all the nodes is not always power efficient and it should be used only when needed. Furthermore, the amount of communication between nodes is not evenly distributed, some nodes communicate more frequently than others, similar observations for both SPLASH and Parsec are presented by Barrow-Williams, et al. [3].

These observations motivate us to explore alternative mNoC designs. Ideally, we want to match NoC power consumption to application communication patterns. Specifically, *destinations with more frequent communication should use a mode that consumes less power, while destinations with less frequent communication should use modes that consume more power.* Importantly, all the modes incur the same latency. The remainder of this paper describes our initial steps toward achieving this objective.

### 3. mNoC Power Topologies

This section describes the foundational concepts that enable mNoC designs that achieve the performance of a full crossbar with energy savings approaching low-radix clustered designs. The key insight is to create a crossbar structure where connectivity (reachable destination nodes from a given source) is a function of source power—called a *power topology*. **The destinations reachable in a given power mode do not need to be contiguous on the physical waveguide.** To our knowledge, this is the first work to define a communication system for on-chip optical networks where on a given waveguide (and thus fixed physical topology) a source node can communicate with a given set of destinations as a function of the source input power. Our goal in this section is to first define the concept of a power topology and then to introduce the design parameters and methods available for architects to create various power topologies. The next section discusses architecting specific power topologies.

#### 3.1 Power Topology Definition

The intuition behind a power topology is to augment a crossbar structure with a set of power modes where different sets of destination nodes are reachable in each power mode.

Broadcast is the highest power mode and can reach all possible destinations, additional lower power modes can only reach decreasing subsets of destinations. Importantly, destinations in a low power mode are also reachable in all higher power modes, and power modes do not need to be uniform across all sources.

More formally, the global power topology  $PT$  of an  $N \times N$  crossbar is the union of each source’s local power topology  $PT = \bigcup_{n=0}^{N-1} PT_n$ . A local power topology  $PT_n$  is an ordered set of  $M_n$  power modes with  $P_{mode_{n,i}}$  Watts, such that  $P_{mode_{n,i}} < P_{mode_{n,j}}, \forall i < j$ , where  $i, j \in \{0..M_n - 1\}$ , and the set of destination nodes reachable in each power mode ( $M_{dest_{n,i}} \subseteq \{0..N - 1\}$ ). The destination nodes reachable in a given power mode are also reachable in all higher power modes ( $M_{dest_{n,i}} \subseteq M_{dest_{n,j}}, \forall i < j$ ). The highest power mode contains all destinations ( $M_{dest_{n,M_n}} = \{0..N - 1\}$ ) and requires the maximum source power. Specifying the number of modes and which nodes are in each mode of the power topology is the architect’s job, as described in Section 4. The remainder of this section describes how to implement a given power topology, and obtain the values of  $P_{mode_{n,i}}$  to minimize the total power for a given topology.


#### 3.2 Implementing Power Topologies

In this paper we focus on the SWMR crossbar structure, where each source node has its own dedicated waveguide(s); however, the approach is general and could be applied to other photonic crossbar structures. Without loss of generality, we consider a single waveguide with a single source and  $N-1$  destinations and describe how to create different power modes. For clarity we omit the subscript  $n$ .

Different power modes are enabled by two key aspects of mNoC technology: 1) waveguide design and 2) current controlled QD\_LED emission (any current controlled light source could be used). Our power topologies are static since they are defined at design time. As part of our ongoing research we are exploring the feasibility of dynamic power topologies. Below we elaborate on the two design aspects for designing static power topologies.

##### 3.2.1 Waveguide Design

Our goal is to implement a power topology to achieve the lowest overall power consumption by minimizing the total source power. Recall, a source QD\_LED injects power into the waveguide, which is dissipated due to waveguide loss, insertion loss incurred by various optical devices (usually fixed), and splitters that divert a fraction of the power to each destination’s receiver. The last factor is the key parameter for creating power topologies and is dependent on waveguide splitter design. Below we outline how to compute appropriate splitter parameters for a given power topology.



**Figure 4.** Source Power Model

Assume we are provided a local power topology (i.e., the number of power modes and the set of destination nodes in each power mode) and an expected fraction of overall communication in each mode  $w_m$ . The total source power is the sum of the power for each mode weighted by the fraction of communication in each mode (see Equation 1).

$$P_{src} = \sum_{m=0}^{M-1} w_m P_{mode_m} \quad (1)$$

Equation 2 represents the relationship between source power and waveguide losses (including splitters) for an arbitrary source injected power  $P_i$  at source  $i$  (See one illustration in Fig. 4).

$$P_i = \sum_{j=0, j \neq i}^{N-1} \frac{\beta_j P_{min}}{L^{|j-i|} S_j (0.5(1+\theta) - S_i \theta) \prod_{\substack{k=\min(i,j), \\ k \neq i,j}}^{\max(i,j)} (1 - S_k)} \quad (2)$$

The sum is over all destinations on the waveguide. The numerator in Equation 2 represents the amount of power incident on  $j$ 's receiver in terms of the minimum required power  $P_{min}$  (which considers the insertion loss of various optical devices and photoreceiver mIOP) where in general  $\beta_j \geq 0$ . In the denominator,  $L^{|j-i|}$  is the power loss along the waveguide between the source  $i$  and the destination  $j$ .  $S_j$  is the fraction of power diverted by the destination's splitter. The third term in the denominator represents how the source power splits ( $S_i$ ) into two directions ( $\theta \in \{1, -1\}$  identifies the direction relative to the source node  $i$ ). The product term accounts for the power diverted by all splitters between  $i$  and  $j$ .

For a single mode power topology (broadcast) total power is minimized if  $\beta_j = 1, \forall j$ . However,  $\beta_j$  is not a directly available design knob, it simply reflects the relative amount of  $P_{min}$  received by the destination, which depends on the distance from the source and the splitter design. If node layout is fixed along the waveguide, then splitter design ( $S_j$ ) is the only remaining design option available to ensure  $\beta_j = 1$  and to minimize total power.

When we consider multiple power modes, the goal is to minimize total power,  $P_{src}$  in Equation 1, not minimize power for a single mode. To achieve this we seek to design splitters,  $S_j$ , that minimize  $P_{src}$ . The intuition behind our

design process is that for a given power mode  $m$  all nodes reachable using power  $P_{mode_m}$  must receive at least  $P_{min}$  (i.e.,  $\beta_j \geq 1$ ) and that the nodes unique to that mode receive exactly  $P_{min}$  (i.e.,  $\beta_j = 1$ ). This basic observation allows us to set up a system of equations (one for each mode) that can be used to solve computationally for  $S_j$ 's that minimize  $P_{src}$ . Appendix A provides further details on the splitter design process.

The primary constraint in our designs is that destinations in low power modes are also reachable in higher power modes. It is important to note that nodes in a given power mode can be non-contiguous; that is nodes reachable only in a higher power mode may be physically closer to the source than nodes reachable in lower power modes. The intuition is that the high power node's splitter diverts a very small fraction of the power to its receiver, therefore a much higher power is required to reach  $P_{min}$  at the receiver. In contrast, the low power node's splitters divert a larger fraction of power. This unique capability enables the creation of arbitrary power topologies.

### 3.2.2 Implementing Source Power Modes

Similar to micro-ring resonators, the QD\_LED can be driven by an integrated high-speed current driver. The QD\_LED output intensity can be adjusted either through pulse-width modulation or dynamic gain adjustment of the driver. Since the required output power (per source-destination pair) is static, software can store a table of constants for each power mode and augment packet transmission with control bits which set the QD\_LED output power. This same table can also store the mapping of logical thread IDs to physical cores, or vice versa. On the receiving side, when the input power is below mIOP, especially in low power modes, the input should be treated as noise. Therefore, to reduce the bit error rate (BER), a simple threshold circuit [4] can be used.

## 4. Architecting Power Topologies

The architecture of an mNoC crossbar specifies the global crossbar power topology: the number of power modes and the set of destination nodes in each power mode, for each source node. Power topologies are created on top of an underlying physical topology; in our case an mNoC SWMR crossbar with a serpentine layout. Designing an optimal power topology requires a global optimization which is NP-hard. Therefore, we decompose the problem into a series of steps that enables us to explore progressively more sophisticated designs. For simplicity, we limit our investigations to power topologies where each source has the same number of power modes ( $M_n = M, \forall n$ ). We begin by mapping conventional topologies onto a power topology. This is followed by power topologies that group destination nodes based on distance from the source. We then introduce communication-aware topologies that account for locality of communication between nodes.

7	2	2	2	2	1	1	1	-
6	2	2	2	2	1	1	-	1
5	2	2	2	2	1	-	1	1
4	2	2	2	2	-	1	1	1
3	1	1	1	-	2	2	2	2
2	1	1	-	1	2	2	2	2
1	1	-	1	1	2	2	2	2
0	-	1	1	1	2	2	2	2
	0	1	2	3	4	5	6	7

(a) Clustered Power Topology

7	4	3	3	2	2	1	1	-
6	4	3	3	2	2	1	-	1
5	4	3	3	2	1	-	1	2
4	4	3	2	1	-	1	2	3
3	3	2	1	-	1	2	3	4
2	2	1	-	1	2	3	3	4
1	1	-	1	2	2	3	3	4
0	-	1	1	2	2	3	3	4
	0	1	2	3	4	5	6	7

(b) Distance-Based Power Topology

**Figure 5.** Example Power Topologies: Each row represents the waveguide(s) for the source in a SWMR crossbar mNoC. Numbers indicate the power mode required to reach a destination in a given column (lower numbers = lower power mode).

#### 4.1 Conventional Topologies

A simple approach is to map known topologies into a power topology. One way to accomplish this is to set the number of power modes based on the diameter of the conventional network. For each source, destination nodes are assigned to power modes based on the number of hops in the shortest path from the source.

To visualize a power topology we use an adjacency matrix where each entry indicates the power mode for a source (row) to reach a specific destination (column). Each row represents the waveguide(s) for the source in a SWMR crossbar mNoC. We note that different sources may have very different values for their respective power modes. For example, power mode 2 for some sources may be much higher power than power mode 2 for other sources due to the maximum propagation distance along the waveguide.

Figure 5a shows an 8-node clustered power topology with four nodes per cluster. This corresponds loosely to the clustered topology used in rNoC and c\_mNoC as described in Section 2. In this power topology there are two power modes: a low power mode (1) for nodes within the source’s

cluster and a high power mode (2) with nodes outside the source’s cluster. Each source has three destinations in its lowest power mode, corresponding to the local communication in the conventional physical topology. The remaining destinations are reachable only in the high power mode. For the 256-node rNoC or c\_NoC systems, there are 252 nodes in the high power mode.

This approach can be used to map any known topology (e.g., trees, binary n-cubes, etc.) into a power topology. However, these architectures may not produce the lowest overall power due to a mismatch between the power characteristics of the waveguides and the defined power topology. For example, nodes three and four in Figure 5a are physically close on the waveguide, yet any communication between them requires the high power mode.

#### 4.2 Distance-Based Topologies


An alternative to naively mapping conventional physical topologies into power topologies is to design new power topologies. One simple approach is to specify a power topology based on the  $n$  nearest destinations. Figure 5b shows a 4-mode power topology based on the two nearest destinations. In this design all sources’ local power topologies have the same value of  $n = 2$ . This is not a requirement in architecting power topologies, and we could explore using different distance parameters per source to create a global power topology with unique local power topologies.

Distance-based topologies are well-matched to the underlying power characteristics of the waveguide and capture localized communication between consecutive neighboring nodes on a line. Unfortunately, they may be a mismatch for communication patterns with higher dimensional locality (non-consecutive nodes on a line).

#### 4.3 Communication-Aware Topologies

To achieve the full benefit of power topologies, we’d like to assign destinations with the highest traffic to the low power modes. Achieving this requires more than an architect’s intuition about traffic patterns, but instead must be data driven through profiling. By utilizing communication patterns and frequency we can potentially identify the optimal number of modes and the assignment of destinations to each mode. Nonetheless, it may be difficult to create a single static power topology that is the best for all benchmarks.

Our goal is to find the power topology that provides the lowest power, which could be achieved by searching over the entire power topology space. In the extreme case, a power topology could have a dedicated mode for each destination. However, finding an optimal power topology has complexity  $O(M^{N-1})$ , which is intractable for large systems with more than a few modes (i.e.,  $N = 256$ ). In this case, heuristic methods must be utilized, and may result in suboptimal power topologies. Developing appropriate algorithmic techniques to fully explore the power topology space is beyond the scope of this paper.



**Figure 6.** mNoC Single-Mode Power Profile


We explore communication-aware power topologies with two and four modes. In general, we apply the “*more is less, less is more*” philosophy and sort the destinations by communication frequency with the source. For the two-mode case, we iterate over all possible binary partitions of the sorted destinations into the two power modes to create  $N - 2$  two mode power topologies. Starting with only the highest frequency destination in the low power mode, we then progressively move destinations to the low power mode until only the lowest frequency destination is in the high power mode. We then choose the power topology with the lowest overall power.

For the four-mode case we apply various *ad hoc* heuristics to assign destinations from the sorted list to the four power modes. We consider several partitions of the sorted destinations into the four modes, such as  $\{64,64,64,63\}$ ,  $\{1,1,2,251\}$ ,  $\{4,120,53,78\}$ , but found the latter to be best, based on manual greedy assignment. We leave a methodical exploration of this space to future work.


Application specific power topologies can be created by applying the above methods to a single application. This could be important for embedded or high-performance ASIC designs where there are limited and well-known usage scenarios. We also note that even the distance-based power topologies could potentially benefit from more accurate information on communication frequency since it might impact splitter design, as discussed previously.

#### 4.4 Thread Mapping

The base mNoC power topology operates by having each source broadcast a packet on its dedicated waveguide(s) and each node locally determine if it is the destination. This single mode topology provides the maximum connectivity: each source can reach all destinations. However, even in this single-mode topology, different source nodes can require different power. Consider the serpentine waveguide layout where all waveguides terminate in the same regions of the chip. For this layout, source nodes occupy locations on their respective waveguide that range from either end of the waveguide to the middle of the waveguide. The nodes in the middle do not require as much source power as the end



(a) Naive Mapping (b) QAP Mapping  
Color indicates communication intensity (dark = high, light = low)



(c) Naive 2-mode Power Topology (d) QAP 2-mode Power Topology  
Color indicates power mode (dark = low, light = high). Each row represents the waveguide(s) for a source node in a SWMR crossbar. Note the non-contiguous destinations in each power mode.

**Figure 7.** Thread Mapping and Power Topologies (water\_spatial). Rows are sources and columns are destinations.

nodes since their signals propagate only half the distance as the end node signals.

The variation in required source power creates a power profile for the overall crossbar. Although each source has a single fixed power, the power for each source depends on its location (see Figure 6). Given this power profile and the observation of workload communication locality, it is beneficial to map frequently communicating threads to the cores located near the center of the waveguide.

Even for different layouts where waveguides terminate in different regions of the chip, a power profile still exists. Cores are arranged in a 2D plane, even a more complex layout of waveguides cannot create a uniform distance from each source to each destination.

Thread mapping can be achieved either offline or online if the workload runs long enough to warrant migration. Our mapping problem is an instance of the quadratic assignment problem (QAP), which is NP-hard, therefore we use heuristic methods. We explore both Taboo [32] and simulated annealing [10], and find that Taboo generally performs best.

#### 4.5 Discussion

Figure 7 shows the communication pattern for the water\_spatial benchmark before (Figure 7a) and after running Taboo (Figure 7b). Each row is a source node (thread) and each column is a destination node, darker colors represent a larger amount of communication. In contrast to the naive

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 2.** Simulation configuration

mapping, the new mapping has very high density communication clustered around the middle nodes, which require lower overall power than the nodes near the ends of the waveguide.

Figure 7 also shows the mode assignment for a 2-mode power topology specific to *water\_spatial* before (Figure 7c) and after running Taboo (Figure 7d). Each row represents a source node’s waveguide(s) and each column is a destination node, a dark color indicates a specific destination node (x axis) is in the low power mode of a source node (y axis). We observe that this power topology matches the communication patterns from Figure 7a and 7b, with more frequently communicating nodes in the low power mode. These figures also show the non-uniform nature of communication-aware power topologies: *source nodes have different local power topologies, and the general framework does not require consecutive destinations within a given power mode.*

In this paper we perform thread mapping based on the single mode power topology and thus the assignment accounts for only the waveguide loss between a source and destination. A more general approach would perform a joint optimization of power topology design and thread mapping. We leave exploring additional heuristic techniques to perform this even more complex assignment as future research.

## 5. Evaluation

In this section, we evaluate the power consumption of mNoC for various power topologies, using naive and QAP thread mapping. Our results show that power topologies combined with QAP thread mapping achieves the goal of performance equal to a full crossbar, but with power consumption close to the smaller clustered topology. Thread mapping is critical for achieving this goal, where two power modes capture a significant fraction of the opportunity (46% reduction), but a 4-mode power topology can provide an additional 5% reduction. We first present our evaluation methodology and then discuss the impact on power consumption of several different design options. The options include naive distance-based power topology, thread mapping, communication-aware power topologies, and application-specific power topologies.

QD_LED energy efficiency	10%
QD_LED 1-to-0 ratio	1
Waveguide loss	1dB/cm
Coupler loss	1dB
Photodetector mIOP	10 $\mu$ W
Power loss of chromophores	5 $\mu$ W for 10 $\mu$ W mIOP
Optical splitter	0.2dB

**Table 3.** Optical energy parameters

### 5.1 Methodology

We create our mNoC topology in the Graphite [25] simulator and run all the simulations in full simulator mode. We explore 256-core systems and the simulation configuration is summarized in Table 2. We use the MOSI directory-based cache coherence protocol provided in Graphite along with its built-in models for contention in the system. The total O/E and E/O latency is about 200 ps and is modeled as 1 cycle in the nanophotonic link traversal time. We assume a die size of 400mm<sup>2</sup>, therefore the waveguide’s total length is approximately 18cm. We conservatively assume the speed of light in the waveguide is about 10cm/ns, which means 1.8ns to travel the longest distance, corresponding to a worst case of 9 cycles for a 5GHz clock. All electrical links are modeled as 1 cycle [8] for the rNoC and clustered mNoC (c-mNoC).

We run 12 benchmarks from the SPLASH benchmark suite [37] to evaluate performance. As mentioned in Section 2, a SWMR single-mode mNoC crossbar improves the performance by 10% compared with rNoC, while c-mNoC has similar performance to rNoC. To evaluate power consumption, we obtain traces of communication packets from all 12 benchmarks executing on mNoC and build a power model to explore the impact of power topologies and thread mapping on power consumption. The key parameters used in the power model are shown in Table 3. The baseline for comparison is total network power for a 256 × 256 single mode mNoC with naive thread mapping, shown in Table 4. For rNoC we calculate power using similar methodology to Pang et al. [30]; however we bias in favor of rNoC by using 0.1 $\mu$ W for photodetector mIOP vs. the 10 $\mu$ W mIOP for mNoC as described in 2.2. We further bias results toward rNoC by assuming 10% QD\_LED efficiency instead of the 18% assumed in [30]. The electrical buffer power is determined using models described by others [19, 27, 28] which may slightly bias in favor of mNoC due to our increased number of buffers.

Our baseline results reveal that the clustered rNoC (radix-64 optical crossbar) consumes 36W, with 23W in ring trimming and a 5W laser source, lower than previous results [27, 36] for SWMR optical crossbars (see Figure 10 in Section 5.7 for a complete breakdown). The average power for the base mNoC is 20.94W, lower than the clustered rNoC. Furthermore, mNoC is energy proportional, applications with higher network utilization (e.g., radix) require



Benchmark	Power (W)	Benchmark	Power (W)
barnes	7.05	water_s	5.28
radix	120.34	water_ns	6.08
ocean_c	12.31	cholesky	5.14
ocean_nc	24.23	lu_cb	7.79
raytrace	3.99	lu_ncb	43.70
fft	11.41	volrend	3.99
average	20.94		

**Table 4.** Base mNoC Power Consumption

Symbol	Definition
M	Mode
T	Thread mapping with QAP
N	Naive distance-based mode assignment
G	General mode assignment based on sampled weights
C	Custom power topology design
U	Uniform traffic pattern for splitter design
W	Weighted traffic pattern for splitter design
S	Sampled traffic weights for splitter design

**Table 5.** Design Symbols

high power ( $> 100W$ ); however we show that power aware thread mapping and power topologies reduce this below rNoC, to  $\sim 20W$ , while maintaining the performance benefits of high-radix crossbars. We use the symbols shown in Table 5 to create a notation for the various power topologies in our figures; for example, 2M\_T\_N.S4 means 2 modes, QAP thread mapping, naive distance-based, and splitter design based on weights from sampling 4 applications.

## 5.2 Naive Distance-Based Power Topologies

We evaluate two distance-based power topologies: 1) a 2-mode topology with the 128 closest destinations assigned to the low power mode of a source, and the remaining destinations assigned to the higher power mode, and 2) a 4-mode power topology based on groups of 64 nearest nodes. To calculate overall source power (Equation 1), and the corresponding splitter design, we explored several communication weightings (e.g., 66% / 33%, uniform, etc.) and found qualitatively similar results across all weightings (see Section 5.6). For clarity we only present results for uniform weights. Figure 8 shows the power consumption of the two distance-based power topologies with naive and QAP thread mapping normalized to the single-mode mNoC with naive thread mapping (1M). Here we consider only the naive thread mapping, we discuss QAP thread mapping below. The 4-mode distance-based power topology reduces the total mNoC power by 12% while the 2-mode distance-based power topology reduces power by 10% on average (harmonic mean). Ocean\_nc and radix show the largest benefits of power topologies with reductions 15%-20%. We also implement 256-node clustered 2-mode power topology similar

to Fig. 5a with naive thread mapping; however it only reduces mNoC power by 1% on average, demonstrating that distance-based power topologies are superior to clustered power topologies.

## 5.3 Impact of Thread Mapping

Clustering communication so that frequently communicating threads are allocated to cores near the center of the waveguide provides greater opportunity to exploit the different power modes in a given power topology. Figure 8 supports this hypothesis; QAP thread mapping helps all three power topologies. The average power reductions over the base mNoC are 27%, 38%, and 39% for the 1-mode, 2-mode, and 4-mode power topologies, respectively. Ocean\_nc and radix show the most improvement from QAP thread mapping with power topologies providing as low as 26%-32% of total network power vs. the base single-mode mNoC power topology. Although QAP thread mapping has larger impact on total mNoC power than power topologies alone, the best overall design uses a 4-mode power topology along with QAP thread mapping.

## 5.4 Communication-Aware Mode Assignment

Communication-aware power mode assignment follows the “*more is less, less is more*” philosophy by utilizing communication information from benchmarks. We first sort all the destination nodes according the communication frequency with a given source node and then assign more frequent destinations to lower power mode and less frequent destinations to higher power modes. In this section we explore two methods for obtaining communication frequency: sampling from four benchmarks (lu\_cb, radix, raytrace, water\_s) and averaging across all 12 benchmarks. A more statistically significant exploration of sampling, while interesting, is beyond the scope of this paper.

To isolate the impact of communication-aware mode assignment, we compare to a distance-based power topology that uses the sampled traffic for splitter design. Therefore, the difference in total mNoC power is mainly due to different mode assignment. As shown in Fig. 9, communication-aware mode assignment generally works better than naive distance-based mode assignment. Using weights from the average of all 12 benchmarks reduces power by 7% compared with the distance-based mode assignment for the 2-mode design. The reduction is increased to 10% for the 4-mode design. Moreover, communication-aware mode assignment based on the average of 12 benchmarks works better than the one based on 4 sampled benchmarks. As expected, when more communication information is available, the better the design.

A 4-mode power topology produces qualitatively similar results, but quantitatively is the best overall design when sampling across all 12 benchmarks with 49% of the base mNoC’s power on average vs. 53% for the 2-mode power topology. Compared to the best distance-based power topol-






Figure 8. Distance-Based Power Topology With and Without Thread Mapping



(a) Two-Mode Power Topologies



(b) Four-Mode Power Topologies

Figure 9. Communication-Aware Mode Designs vs Distance-Based Mode Designs

ogy that uses the uniform traffic for splitter design in Fig. 8, the best 4-mode power topology with communication-aware mode assignment reduces the total mNoC power by 12%. This demonstrates the benefit of communication aware power topologies.

### 5.5 Application Specific Designs


To understand how well the various general power topologies work, we create application specific power topologies tailored to each benchmark. We examine the source power consumption for the application specific designs with QAP thread mapping and naive mapping. They both have similar characteristics, but do not have a significant improvement over naive distance-based power topologies (about 8%). The “keep it simple” design philosophy may apply. However, for embedded systems or situations with known specific com-

munication patterns, a custom power topology may be beneficial.

### 5.6 Splitter Design Sensitivity Analysis

As explained in Section 3.2 and Section 4.3, the communication patterns of benchmarks could potentially affect the splitter design. We evaluated the sensitivity of splitter design to communication traffic weights using the application specific 2-mode power topology with QAP thread mapping as the basic design. The different communication weights we consider are: uniform, 66%/33%, 33%/66%, samples from 4 benchmarks (S4), and the average across all 12 benchmarks (S12).

Our results reveal minimal variation in power reduction across the different weights (within 2%), but all produce over a 40% reduction in total mNoC power on average. One potential explanation for this limited impact of weights on



**Figure 10.** Total NoC Energy Consumption

splitter designs is that the total source power for given a source core is a function of both weights and splitter ratios. Changes in the weights can be compensated by changes in the splitter ratios resulting in similar overall power. Therefore, the traffic pattern, as it relates to splitter design, has minimal impact on overall power consumption.

### 5.7 Overall mNoC Energy Consumption

As mentioned in Section 2.2, our goal is to design a single physical network with the same performance as the high-radix mNoC, but with the power savings of the lower-radix clustered topology (c\_mNoC). Figure 10 shows the total energy normalized to rNoC for baseline (mNoC), c\_mNoC, and the best power topology mNoC (PT\_mNoC=4M.T.G.S12) with static splitters. The power model for rNoC is similar to the one proposed by Joshi et al. [19]. We use  $20\mu W/\text{ring}$  over 20K temperature range as thermal tuning power to favor rNoC. More accurate ring models [26, 39] indicate much higher ring tuning power. Since the dominant power for rNoC is ring tuning instead of O/E, it shows little or no benefit to increase mIOP for rNoC (see Section 2.2). Therefore, we keep rNoC’s photodetector mIOP as  $1\mu W$  instead of  $10\mu W$  used by all mNoC related topologies. From these results we see that the best power topology with thread mapping requires only 28% of rNoC’s energy, close to the 21% required by c\_mNoC and much lower than the base mNoC (1M power topology) of 57%. Note that c\_mNoC has smaller source and OE/EO power than mNoC, because the lower radix requires shorter waveguide length and fewer optical devices.

Our power topology has at least the same performance as the base radix-256 mNoC (1M). Performance could be improved by considering QAP thread mapping, because it reduces the distance between source and destination, and potentially reduces latency. Thread mapping will have minimal impact on rNoC’s power consumption since it is dominated by thermal ring tuning. Similarly, c\_mNoC’s power is dominated by electrical components. There may be some impact on performance of the clustered topologies (rNoC and c\_mNoC) by thread mapping, however only three nodes are near neighbors and communication with all others re-

quires traversal of two local routers and the larger crossbar. Therefore, we don’t expect to observe too much benefit from thread mapping, nonetheless, we are currently exploring this design point.

## 6. Related Work

We discuss related work based on three broad categories: current nanophotonic NoC designs, power optimization in both on-chip networks and wireless networks, and thread mapping methods to optimize the network in terms of energy, performance or congestion in NoCs. Although different power optimization methods have been proposed for on-chip networks, the concepts proposed in this paper such as logical power mode and power topology in nanophotonic NoCs are new.

Bus-based crossbars are very popular in the current nanophotonic NoCs. A typical nanophotonic crossbar is usually designed using one of the following structures: Multi-Write-Single-Read (MWSR) as proposed for Corona [36], Single-Write-Multi-Read (SWMR) shared-bus crossbar proposed by Kirman et al. [21] and Pan et al. [28], or Multi-Write-Multi-Read (MWMR) which combines the above two and proposes a reduced number of channels design as proposed for Flexishare [27]. However, due to the power inefficiency from the off-chip laser source and thermal ring tuning power, and the ring resonator’s nonlinearity, these designs are difficult to scale to more than radix-64. Nitta et al. [26] presents a thermal model to calculate the ring tuning power, which is much bigger than the heating power calculated with the popular per ring fixed cost method. Even with further optimization techniques [39], the ring power consumption is still a dominant factor. Other topologies based on ring resonator switches are also explored, such as a 2D grid NoC from Phastlane [9] and a torus-based hierarchical hybrid NoC called THOE [41]. Xue et al. [40] proposes a freespace optical interconnect that uses VCSELs as light sources.

As power becomes a larger concern in current computing systems, various power optimization methods are proposed for on-chip networks. Recent work from Koka et al. [22] presents a power constrained method for designing and evaluating the nanophotonic interconnects. Their results show that under reasonable device assumptions for a given input optical power, point-to-point networks have better power and performance characteristics than switched-based networks. Catnap [11] proposes a power proportional NoC design which divides a single NoC into multiple subnetworks to exploit the benefits of power gating. We could apply this same method on mNoC by deactivating waveguides per source to decrease bandwidth and reduce power. Other work explores different physical topologies to reduce NoC energy and/or latency [12, 17, 18]. Finally, any transmission line or wireless communication system can utilize the same prin-

ciples of power modes based on distance to destinations, e.g., [2, 7].

Thread mapping is previously used to design energy and/or latency-aware on-chip networks. Hu et al. [16] proposes an efficient branch-and-bound algorithm to map cores to a regular tile-based network to save total energy. Hoefler et al. [15] discusses efficient application communication pattern mapping to sparse network topologies. They demonstrate a heuristic method based on graph similarity to reduce network congestion and also improve application communication performance. Their methods could potentially be used for joint optimization of thread mapping and power topology design.

## 7. Conclusion

Energy efficiency and power dissipation are increasingly important design goals for computer systems. Manycore processors exacerbate this problem since on-chip interconnect (NoC) can become a significant portion of overall power consumption. This paper utilizes a new molecular-scale photonic NoC technology (mNoC) and shows how to architect power efficient designs through the use of power topologies. A power topology is a set of power modes and an assignment of destination nodes reachable in these power modes, for each source node in the system. Power topology implementations are enabled by careful waveguide design and the unique mNoC capability of current-controlled on-chip QD\_LEDs for light sources.

The benefits of mNoC power topologies include scaling to high-radix crossbars, to achieve high performance, while providing power savings close to smaller-radix crossbars. We present several different power topologies for a 256-core system with two and four power modes, including distance-based and communication-aware topologies that reduce network power by  $\geq 10\%$  compared to a baseline mNoC with a single power mode that can reach all destinations. Further improvements, additional 40% reduction, are possible with intelligent thread mapping that maps frequently communicating threads to cores with lower power communication requirements. We find that distance-based power topologies capture a significant fraction of the overall benefit in mNoC power reduction 40%, but that utilizing application information for mode assignment can achieve over a 50% reduction. Overall, mNoC with power topologies achieves performance 10% higher than a clustered ring-resonator photonic NoC, while reducing energy by 72%.

The work in this paper represents our initial efforts at exploiting emerging molecular-scale photonics to architect novel power efficient on-chip interconnects. The design space is very large, and we've explored only a small portion. Areas of future work include, but are not limited to: dynamic power modes, joint optimization of thread mapping and power topology design, applying power topologies to embedded systems with well-defined IP blocks and com-

munication patterns, and exploring mNoC's ability to multicast/broadcast when used in coherence protocol design.

## Acknowledgements

This material is based upon work supported by, or in part by, the National Science Foundation (CCF-0702434), the Defense Advanced Research Projects Agency and the U. S. Army Research Office under grant number W911NF-13-1-0096 (Distribution Statement A: Approved for Public Release, Distribution Unlimited). The views, opinions, and/or findings contained in this article are those of the author(s) and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government. We thank the anonymous reviewers and the Duke self-assembled systems group for their comments that improved the quality of the paper.

## Appendix A

For a  $N \times N$  crossbar, to solve for  $S_j$  at the destination  $j$  given a source node  $i$ , where  $i, j \in \{0..N-1\}$ , we build on the observation that destination nodes unique to a given power mode  $m \in \{0..M\}$  have the same  $\beta_j$  values which we call  $\gamma_m$ :  $\beta_i = \beta_j = \gamma_m, \forall i, j \in M_{dest_m} - M_{dest_{m-1}}$ , and for each source mode power  $P_{mode_k}, k \in \{0..M\}$  we define unique  $\gamma_{k,m}$ . We then further simplify by defining all  $\gamma_{k,m}$  in terms of  $\gamma_{0,m}$  in power mode 0 to obtain  $\alpha_m = \gamma_{0,m}, m \in \{0..M\}$ , where  $\alpha_m \in (0..1], \alpha_0 = 1$ .

Equations 3-6 show abbreviated versions of Equation 2 for a four mode ( $M = 4$ ) power topology, which connect Equation 2 to Equation 1. The equations are based on an observation that given a fabricated waveguide, a given source power is proportional to *any* destination's received power, and that a change in the source power results in a proportional change in *all* destination's received power.

$$P_{mode_0} = f(P_{min}, \alpha_1 P_{min}, \alpha_2 P_{min}, \alpha_3 P_{min}) \quad (3)$$

$$P_{mode_1} = f\left(\frac{1}{\alpha_1} P_{min}, P_{min}, \frac{\alpha_2}{\alpha_1} P_{min}, \frac{\alpha_3}{\alpha_1} P_{min}\right) \quad (4)$$

$$P_{mode_2} = f\left(\frac{1}{\alpha_2} P_{min}, \frac{\alpha_1}{\alpha_2} P_{min}, P_{min}, \frac{\alpha_3}{\alpha_2} P_{min}\right) \quad (5)$$

$$P_{mode_3} = f\left(\frac{1}{\alpha_3} P_{min}, \frac{\alpha_1}{\alpha_3} P_{min}, \frac{\alpha_2}{\alpha_3} P_{min}, P_{min}\right) \quad (6)$$

We can then rearrange the summand of Equation 2 to express  $S_j$  in terms of  $\alpha_m, P_{min}, \& L$ . We then iterate over all  $\alpha_m$  from 0 to 1 in increments of 0.1 to obtain values for each power mode.

Note that by determining the values that minimize  $P_{mode_0}$ , all higher power modes can be computed by  $P_{mode_i} = \frac{P_{mode_0}}{\alpha_i}$ . Once we have the alpha values we can substitute into the splitter equation to obtain the appropriate splitter values for use in fabrication. Since we have one equation with many unknowns, we solve for alpha values computationally by iterating over all alpha values (better results may be achieved by using steps smaller than 0.1).

## References

- [1] R. Arians, A. Gust, T. Kummell, C. Kruse, S. Zaitsev, G. Bacher, and D. Hommel. Electrically driven single quantum dot emitter operating at room temperature. *Applied Physics Letters*, 93(17):173506–173506, 2008.
- [2] Wonseok Baek, David SL Wei, and C-CJ Kuo. Power-aware topology control for wireless ad-hoc networks. In *Wireless Communications and Networking Conference, 2006. WCNC 2006. IEEE*, volume 1, pages 406–412. IEEE, 2006.
- [3] 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*, pages 86–97. IEEE, 2009.
- [4] Valeriu Beiu, José M Quintana, and María J Avedillo. Vlsi implementations of threshold logic-a comprehensive survey. *Neural Networks, IEEE Transactions on*, 14(5):1217–1243, 2003.
- [5] A. Biberman, K. Preston, G. Hendry, N. Sherwood-Droz, J. Chan, J.S. Levy, M. Lipson, and K. Bergman. Photonic network-on-chip architectures using multilayer deposited silicon materials for high-performance chip multiprocessors. *ACM Journal on Emerging Technologies in Computing Systems (JETC)*, 7(2):7, 2011.
- [6] J. Chan, G. Hendry, K. Bergman, and L.P. Carloni. 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):1507–1520, 2011.
- [7] Benjie Chen, Kyle Jamieson, Hari Balakrishnan, and Robert Morris. Span: An energy-efficient coordination algorithm for topology maintenance in ad hoc wireless networks. In *Proceedings of the 7th annual international conference on Mobile computing and networking*, pages 85–96. ACM, 2001.
- [8] G. Chen, H. Chen, M. Haurylau, N.A. Nelson, D.H. Albonesi, P.M. Fauchet, and E.G. Friedman. Predictions of cmos compatible on-chip optical interconnect. *Integration, the VLSI journal*, 40(4):434–446, 2007.
- [9] M.J. Cianchetti, J.C. Kerekes, and D.H. Albonesi. Phastlane: a rapid transit optical routing network. In *36th annual international symposium on Computer architecture*, pages 441–450, 2009.
- [10] David T Connolly. An improved annealing scheme for the qap. *European Journal of Operational Research*, 46(1):93–100, 1990.
- [11] Reetuparna Das, Satish Narayanasamy, Sudhir K Satpathy, and Ronald G Dreslinski. Catnap: Energy proportional multi-network-on-chip. In *to appear in Proceeding of the 40th annual international symposium on Computer architecture*, 2013.
- [12] Haytham Elmiligi, Ahmed A Morgan, M Watheq El-Kharashi, and Fayez Gebali. Power-aware topology optimization for networks-on-chips. In *Circuits and Systems, 2008. ISCAS 2008. IEEE International Symposium on*, pages 360–363. IEEE, 2008.
- [13] Ashwini Gopal, Kazunori Hoshino, Sunmin Kim, and Xiaojing Zhang. Multi-color colloidal quantum dot based light emitting diodes micropatterned on silicon hole transporting layers. *Nanotechnology*, 20(23):235201, 2009.
- [14] F Hargart, CA Kessler, T Schwarzback, E Koroknay, S Weidenfeld, M Jetter, and P Michler. Electrically driven quantum dot single-photon source at 2 ghz excitation repetition rate with ultra-low emission time jitter. *Applied Physics Letters*, 102(1):011126–011126, 2013.
- [15] Torsten Hoefler and Marc Snir. Generic topology mapping strategies for large-scale parallel architectures. In *Proceedings of the international conference on Supercomputing*, pages 75–84. ACM, 2011.
- [16] Jingcao Hu and Radu Marculescu. Energy-aware mapping for tile-based noc architectures under performance constraints. In *Design Automation Conference, 2003. Proceedings of the ASP-DAC 2003. Asia and South Pacific*, pages 233–239. IEEE, 2003.
- [17] Yuanfang Hu, Hongyu Chen, Yi Zhu, Andrew A Chien, and Chung-Kuan Cheng. Physical synthesis of energy-efficient networks-on-chip through topology exploration and wire style optimization. In *Computer Design: VLSI in Computers and Processors, 2005. ICCD 2005. Proceedings. 2005 IEEE International Conference on*, pages 111–118. IEEE, 2005.
- [18] Yuanfang Hu, Yi Zhu, Hongyu Chen, Ronald Graham, and Chung-Kuan Cheng. Communication latency aware low power noc synthesis. In *Proceedings of the 43rd annual Design Automation Conference*, pages 574–579. ACM, 2006.
- [19] A. Joshi, C. Batten, Y.J. Kwon, S. Beamer, I. Shamim, K. Asanovic, and V. Stojanovic. Silicon-photonic cmos networks for global on-chip communication. In *Proceedings of the 2009 3rd ACM/IEEE International Symposium on Networks-on-Chip*, pages 124–133. IEEE Computer Society, 2009.
- [20] Pawan Kapur. Scaling induced performance challenges/limitations of on-chip metal interconnects and comparisons with optical interconnects. Ph.D. Dissertation, Stanford University, Stanford, California, 2002.
- [21] N. Kirman, M. Kirman, R.K. Dokania, J.F. Martinez, A.B. Apsel, M.A. Watkins, and D.H. Albonesi. Leveraging optical technology in future bus-based chip multiprocessors. In *Proceedings of the 39th Annual IEEE/ACM International Symposium on Microarchitecture*, pages 492–503. IEEE Computer Society, 2006.
- [22] Pranay Koka, Michael O McCracken, Herb Schwetman, Chia-Hsin Owen Chen, Xuezhe Zheng, Ron Ho, Kannan Raj, and Ashok V Krishnamoorthy. A micro-architectural analysis of switched photonic multi-chip interconnects. In *Proceedings of the 39th International Symposium on Computer Architecture*, pages 153–164. IEEE Press, 2012.
- [23] J.R. Lakowicz. *Principles of fluorescence spectroscopy*, volume 1. Springer, 2006.
- [24] Benjamin S Mashford, Matthew Stevenson, Zoran Popovic, Charles Hamilton, Zhaoqun Zhou, Craig Breen, Jonathan Steckel, Vladimir Bulovic, Mounji Bawendi, Seth Coe-Sullivan, et al. High-efficiency quantum-dot light-emitting devices with enhanced charge injection. *Nature Photonics*, 7(5):407–412, 2013.

- [25] J.E. Miller, H. Kasture, G. Kurian, C. Gruenwald, N. Beckmann, C. Celio, J. Eastep, and A. Agarwal. Graphite: A distributed parallel simulator for multicores. In *High Performance Computer Architecture (HPCA), 2010 IEEE 16th International Symposium on*, pages 1–12. IEEE, 2010.
- [26] Christopher Nitta, Matthew Farrens, and Venkatesh Akella. Addressing system-level trimming issues in on-chip nanophotonic networks. In *High Performance Computer Architecture (HPCA), 2011 IEEE 17th International Symposium on*, pages 122–131. IEEE, 2011.
- [27] 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*, pages 1–12. IEEE, 2010.
- [28] Yan Pan, Prabhat Kumar, John Kim, Gokhan Memik, Yu Zhang, and Alok Choudhary. Firefly: Illuminating future network-on-chip with nanophotonics. In *In Proc. of the International Symposium on Computer Architecture*, 2009.
- [29] Jun Pang, Christopher Dwyer, and Alvin R Lebeck. Exploiting emerging technologies for nanoscale photonic networks-on-chip. In *Proceedings of the Sixth International Workshop on Network on Chip Architectures*, pages 53–58. ACM, 2013.
- [30] Jun Pang, Christopher Dwyer, and Alvin R. Lebeck. mNoC: Large nanophotonic network-on-chip crossbars with molecular scale devices. *ACM Journal on Emerging Technologies in Computing Systems (JETC)*, to appear.
- [31] 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. Green light-emitting diodes with self-assembled in-rich ingan quantum dots. *Applied Physics Letters*, 91:133105, 2007.
- [32] E Taillard. Robust taboo search for the quadratic assignment problem. *Parallel computing*, 17(4):443–455, 1991.
- [33] Dmitri V Talapin and Jonathan Steckel. Quantum dot light-emitting devices. *Mrs Bulletin*, 38(09):685–691, 2013.
- [34] Limin Tong, Rafael R Gattass, Jonathan B Ashcom, Sailing He, Jingyi Lou, Mengyan Shen, Iva Maxwell, and Eric Mazur. Subwavelength-diameter silica wires for low-loss optical wave guiding. *Nature*, 426(6968):816–819, 2003.
- [35] Bernard Valeur et al. *Molecular fluorescence: principles and applications*. John Wiley & Sons, 2013.
- [36] Dana Vantrease, Robert Schreiber, Matteo Monchiero, Moray McLaren, Norman P. Jouppi, Marco Fiorentino, Al Davis, Nathan Binkert, Raymond G. Beausoleil, and Jung Ho Ahn. Corona: System implications of emerging nanophotonic technology. In *Proceedings of the 35th Annual International Symposium on Computer Architecture*, pages 153–164, Washington, DC, USA, 2008. IEEE Computer Society.
- [37] Steven Cameron Woo, Moriyoshi Ohara, Evan Torrie, Jaswinder Pal Singh, and Anoop Gupta. The splash-2 programs: characterization and methodological considerations. In *Proceedings of the 22nd annual international symposium on Computer architecture*, pages 24–36. ACM, 1995.
- [38] V. Wood and V. Bulović. Colloidal quantum dot light-emitting devices. *Nano Reviews*, 1(0), 2010.
- [39] Yi Xu, Jun Yang, and Rami Melhem. Tolerating process variations in nanophotonic on-chip networks. In *Proceedings of the 39th International Symposium on Computer Architecture*, pages 142–152. IEEE Press, 2012.
- [40] Jing Xue, Alok Garg, Berkehan Ciftcioglu, Jianyun Hu, Shang Wang, Ioannis Savidis, Manish Jain, Rebecca Berman, Peng Liu, Michael Huang, Hui Wu, Eby Friedman, Gary Wicks, and Duncan Moore. An intra-chip free-space optical interconnect. In *Proceedings of the 37th Annual International Symposium on Computer Architecture*, ISCA '10, pages 94–105, New York, NY, USA, 2010. ACM.
- [41] Y. Ye, J. Xu, X. Wu, W. Zhang, W. Liu, and M. Nikdast. 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):5, 2012.