

QuT: A Low-Power Optical Network-on-Chip

Parisa Khadem Hamedani*, Natalie Enright Jerger* and Shaahin Hessabi†

*The Edward S. Rogers Sr. Department of Electrical and Computer Engineering, University of Toronto, Toronto, Canada

†Department of Computer Engineering, Sharif University of Technology, Tehran, Iran

Email: {parisa, enright}@eecg.toronto.edu, hessabi@sharif.edu

Abstract—To enable the adoption of optical Networks-on-Chip (NoCs) and allow them to scale to large systems, they must be designed to consume less power and energy. Therefore, optical NoCs must use a small number of wavelengths, avoid excessive insertion loss and reduce the number of microring resonators. We propose the Quartern Topology (QuT), a novel low-power all-optical NoC. We also propose a deterministic *wavelength routing* algorithm based on Wavelength Division Multiplexing that allows us to reduce the number of wavelengths and microring resonators in optical routers. The key advantages of QuT network are simplicity and lower power consumption. We compare QuT against three alternative all-optical NoCs: optical Spidergon, λ -router and Corona under different synthetic traffic patterns. QuT demonstrates good scalability with significantly lower power and competitive latency. Our optical topology reduces power by 23%, 86.3% and 52.7% compared with 128-node optical Spidergon, λ -router and Corona, respectively.

I. INTRODUCTION

As we usher in the many-core era, increasing attention is being paid to communication as a potential power and performance bottleneck. Electrical NoCs are an effective solution for current and near-future on-die communication requirements. However, as future many-core systems require higher throughput and operate under increasingly tight power budgets, the continued scalability of electrical NoCs is questionable [1][2]. Integrated silicon-compatible photonic technology is an alternative approach to communication, which has the potential to support large-scale networks with low power, low latency and high bandwidth [3][4]. Photonics reduce network power consumption by omitting electrical buffering and switching.

Electrical NoCs. Power and latency constraints in future CMPs severely limit the scalability of electrical NoCs. For an electrical mesh with 168-bit flits and 4-flit buffers per input port, the energy to transmit one flit across a 1 mm link and subsequent router (energy-per-flit-per-hop) is 197 pJ [2][5]. The total energy consumed in a clock cycle is proportional to the average number of flits traversing links per clock cycle [2]. Therefore, the power consumption of 128-node mesh is 82 W¹, in 32nm at 3 GHz under uniform traffic, with average link utilization of 0.3 and XY routing. This power consumption is unacceptably high. Evaluation of a 256-node 2D mesh in 22nm shows that the network channel power and buffering power of an electrical NoC exceeds the allocated network power budget by an order of magnitude [1].

In terms of performance, assuming 3 clock cycles per router and a single cycle per link, the average zero-load latency under uniform random traffic is 32 and 43 clock cycles in 128- and 256-nodes mesh NoC, respectively. Zero-load latency does not consider router congestion which has a negative effect

on latency. As the network size grows, the latency in an electrical NoC quickly increases [6]. Since future CMPs need a low power and low latency communication fabric, studying alternatives such as Optical NoCs which can overcome these limitations is valuable.

Optical NoCs. Although nano-photonics promises low-power, high-throughput communication, it presents a number of design challenges. Despite the fact that research in nano-photonics devices and materials is a hot area, there are some technical hurdles that could limit the deployment of optical networks. To overcome these impediments, scalable optical networks should carefully consider power, the number of microrings and the number of wavelengths.

- **Insertion Loss:** To fully realize the potential of nano-photonics, the network must be carefully designed with low power in mind. Laser power is one of the main factors in the power consumption of an optical NoC. The laser power is calculated based on the maximum insertion loss (IL) in an optical system. Therefore, to reduce optical power, we must decrease IL.
- **Number of Microrings:** In terms of reliability, microrings are a major source of faults in an optical system, since they are sensitive to process and temperature variation. An optical NoC should have simple switches to reduce the number of microrings in the path of the optical streams.
- **Number of Wavelengths:** The number of wavelengths used to modulate the data in an optical system, should be reduced to have a low power and practical optical NoC.
 - Wavelength-division multiplexing (WDM) can substantially improve the optical network bandwidth. WDM allows the transmission of different data streams through a single optical waveguide [3]. Although, existing work relies heavily on Dense WDM, the number of wavelengths that can be transmitted in a single waveguide is limited. Current research estimates that less than 100 wavelengths can be supported per waveguide [8].
 - The laser must produce sufficient output power to propagate an optical stream with low bit error rate. The output power of today's multi-wavelength lasers is low [9][10]. Low-power means there is not enough power to modulate data correctly across many wavelengths. As a result, a single-wavelength laser must be used; however, a large number of single-wavelength lasers cannot be attached to the chip. Therefore, the number of wavelengths should be small to reduce the number of lasers.
 - Laser power is proportional to the number of wavelengths in an optical system. Supporting a large number of wavelengths increases the required laser power.

To design a low power optical architecture, we must address all of these challenges. Focusing on one or two of them is not

This research was supported by the Natural Sciences and Engineering Research Council of Canada and the University of Toronto.

¹Power consumption is estimated using the method by Eisley and Peh [7].

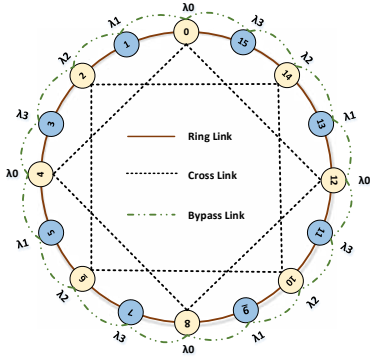


Figure 1: 16-node QuT structure and wavelength assignment

sufficient to have a low power optical NoC. We propose the Quartern Topology (QuT), a novel low-power all-optical NoC targeting key optical challenges. We use passive microring resonators [11] to route optical streams based on their wavelength. QuT is a ring-like topology with strategically placed extra links to reduce the diameter and number of wavelengths required. Our proposed deterministic routing algorithm avoids collisions between streams that need to transmit on the same wavelength.

Contributions. We make the following key contributions:

- A new all-optical architecture: QuT takes steps towards addressing the optical challenges including IL, number of wavelengths and microrings. Therefore, QuT consumes less power compared with state-of-the-art proposals;
- A new deterministic wavelength routing that provides contention-free network traversal while reducing the number of wavelengths required by optimizing the optical switches;
- 128-node QuT achieves power reductions of up to 23%, 86.3% and 52.7% over Spidergon, λ -router and Corona.

II. QUARTERN ARCHITECTURE

In this section, we present our low-power all-optical NoC that features a small number of wavelengths and simple switches that leads to low resource requirements. These features are essential to realizing optical networks in practice and scaling them to large systems. QuT is a contention-free topology; data streams intended for different destinations do not block each other. However, a given destination can only accept traffic from a single source at a time; as a result, QuT requires a control network to serialize access by multiple sources to a particular destination. Both the control and data networks are optical. We first explain the data network and optical switches. Then, we explain the WDM routing and clarify the details of QuT through an example. Finally, we explain the control network.

A. Data Network

QuT uses passive microring resonators, which route optical streams based on their wavelengths. Figure 1 shows a 16-node QuT network. Extending this structure to larger numbers of nodes is straightforward; the connectivity for an arbitrary N nodes is given below. QuT includes two different types of switches and three different connections between nodes which enable a WDM routing approach to communicate between nodes. The link for sending data is chosen based on the distance between the *current* node and the destination. Data traverses *Ring* links when the distance between the current node and the destination is less than $N/4$. Otherwise, *Cross* or *Bypass* links are used for sending data streams. Utilizing these different links reduces the network diameter to $N/4 + 1$.

QuT is based on the following connectivity formulation:

- *Ring links (bidirectional)*
 x_i is connected to $\begin{cases} (x_i + 1) \bmod N \\ (x_i - 1) \bmod N \end{cases}$
- *Cross links (bidirectional)*
if i is even, x_i is connected to $\begin{cases} (x_i + N/4) \bmod N \\ (x_i + 3N/4) \bmod N \end{cases}$
- *Bypass links (unidirectional)*
if i is odd, x_i is connected to $\begin{cases} (x_i + 1) \bmod N \\ (x_i - 1) \bmod N \end{cases}$

Where switch x_i is connected to node i and N is the number of nodes in QuT.

Although bypass links appear to have the same connectivity as ring links, there are two important differences. First, they are unidirectional: they only emanate from odd nodes. Second, their connection within the switch is different from ring links. We use bypass links, instead of additional cross links in odd switches, to prevent the data from being absorbed by the wrong node as we will see in Section II-B2.

B. Router Microarchitecture

In this section, we describe the assignment of wavelengths to nodes and the design of the optical switches.

1) *Wavelength Assignment:* Each node in the QuT data network has a specific, dedicated but not unique wavelength onto which other nodes will modulate their data to communicate with this node. We have designed QuT, its deterministic wavelength routing and its optical switches to use a small number of wavelengths. In an N -node QuT, we use $N/4$ distinct wavelength sets to modulate all data in network. A wavelength set of $\lambda_{(i \bmod N/4)}$ is assigned to node i (Figure 1). Each wavelength set can include one or more wavelengths based on the required network bandwidth.

2) *All-Optical Switches:* QuT requires two different switches, one for even (Figure 2a) and one for odd (Figure 2b) nodes. Employing passive microring resonators in QuT switches eliminates the path reservation phase required in prior work that used electro-optic broadband ring resonators [12]. A source chooses one of the four injection channels in its optical switch, based on the destination's index. Choosing among four injection channels helps reduce the total number of wavelengths required by the network.

Microring resonators in the even and the odd switches are classified in four groups: (a) Add Microring Resonators ($A\mu R$) that multiplex optical streams from injection channels or cross links to ring links (b) Bypass Microring Resonators ($B\mu R$) in the even switches that switch optical streams from bypass links to cross links (c) Cross Microring Resonators ($C\mu R$) in odd switches which turn optical streams from ring links to bypass links (d) Drop Microring Resonators ($D\mu R$) that send the data, modulated on the wavelength of corresponding destination node, from ring links into the switch's detector.

$A\mu R$, $B\mu R$, $C\mu R$ and $D\mu R$ are arrays of microrings, where each microring is sensitive to a specific wavelength. By assuming that $\lambda_{(i \bmod N/4)}$ is devoted to i th node, $D\mu R$ in the i th optical switch is only sensitive to $\lambda_{(i \bmod N/4)}$. $A\mu R$ is used to prepare suitable turns for all optical data streams. Hence, $A\mu R$ is sensitive to all wavelengths in QuT. $B\mu R$ is used to send data on one of the cross links connected to the even node's switch. Therefore, $B\mu R$ is sensitive to all wavelengths in QuT. Even nodes benefit from $C\mu R$ s in odd switches when they want to send data to the destination that is a distance

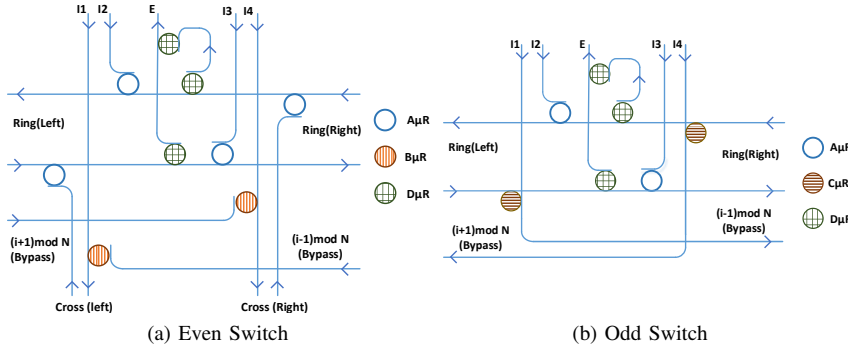


Figure 2: Structure of QuT Switches

of $N/2$ from them. In QuT, nodes with distance $N/4$ or $N/2$ from a source, have the same wavelength because we use $N/4$ different wavelengths. As a result, if a source wants to send data to the destination that is a distance of $N/2$, it should use a path that will bypass the node that is $N/4$ from it to prevent the data stream from being absorbed by the wrong node. Employing $C_{\mu}R$ s in odd switches prepares suitable paths for that purpose. Hence, $C_{\mu}R_L$ ($C_{\mu}R_R$) in the i th optical switch is sensitive to the wavelength assigned to the even switch to the left (right) of the i th optical switch. $C_{\mu}R_L$ is sensitive to $\lambda_{((i+1) \bmod N/4)}$ and $C_{\mu}R_R$ is sensitive to $\lambda_{((i-1) \bmod N/4)}$.

C. QuT WDM Routing

We develop a novel deterministic optical routing algorithm tailored to the QuT architecture and based on wavelength routing. The wavelength assignment coupled with the routing algorithm allows us to realize our contention-free architecture. No two distinct source-destination pairs using the same wavelength will collide in the network; if they use the same wavelength, they will take different routes. Also, using *dedicated deterministic* optical routing helps optimize the optical switches and reduces the number of microrings compared with non-deterministic optical routing algorithm.

First, the source node chooses an injection channel (designated $I\#$) based on the destination. Then, data is modulated onto the destination wavelengths. The optical data stream is transmitted through optical switches and links until it is ejected at the destination. There are several scenarios depending on the source and destination, listed as follows:

- If the source is even and the distance between the source and the destination is:
 - less than $N/4$, I_2 (I_3) is chosen if the destination is to the left (right) of the source. Data is sent on the ring link to the left (right) of the current switch.
 - equal to $N/4$, I_1 (I_4) is chosen if the destination is to the left (right) of the source. Data is sent over the cross link to the left (right) of the current switch. When the distance between two nodes is divisible by $N/4$, these nodes have the same dedicated wavelength.
 - equal to $N/2$, the source and the destination have the same dedicated wavelength. Therefore, I_2 is chosen to send data over the ring link to the left of the current switch. The odd switch to the left of current switch will send the data stream over its bypass link, to bypass the node that is a distance of $N/4$ from the source and prevent the data from being absorbed by wrong node.

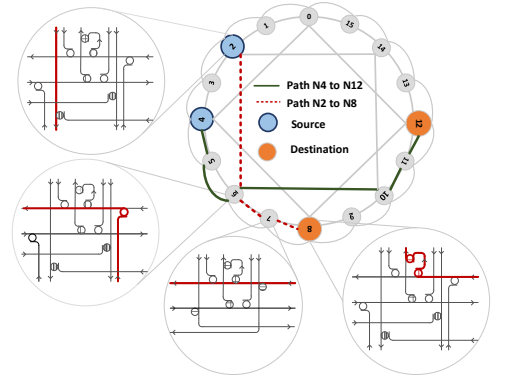


Figure 3: Routing Example: N_4 to N_{12} , N_2 to N_8




Figure 4: Control Network Structure

- greater than $N/4$ and not equal to $N/2$, I_1 (I_4) is chosen if the destination is to the left (right) of the source. The cross link to the left (right) of the current switch is used.
- If the source is odd and the distance between the source and the destination is:
 - less than or equal to $N/4$, I_2 (I_3) is chosen if the destination is to the left (right) of the source. Data is sent on the ring link to the left (right) of the current switch.
 - greater than $N/4$, I_1 (I_4) is chosen if the destination is to the left (right) of the source. Data is sent over the bypass link to the right (left) of the current switch.

D. Example

Figure 3 gives an example of routing in QuT. A common wavelength (λ_0) is used for both data streams from N_4 to N_{12} and N_2 to N_8 ; data transmission through different paths prevents optical collision. If N_2 wants to send data to N_8 , it chooses injection channel I_1 . Therefore, data is transferred over a cross link connected to N_6 . Then, the data stream is turned onto the ring link by $A_{\mu}R$ between the cross and ring links in N_6 . Data is transferred through ring link without changing direction in N_7 . Finally, in N_8 , data turns into the ejection channel through $D_{\mu}R$. If N_4 wants to send data to N_{12} , it chooses injection channel I_2 . Therefore, data is transferred over a ring link connected to N_5 by $A_{\mu}R$ in N_4 . Then the bypass link transmits the data through $C_{\mu}R$ in N_5 . In N_6 , data is transferred through the cross link by $B_{\mu}R$ to N_{10} . In N_{10} , data is sent over the ring link by $A_{\mu}R$. Data is transferred through ring link without changing direction in N_{11} . Finally, in N_{12} , data is ejected by $D_{\mu}R$.

E. Control Network

QuT checks the status of the destination's ejection channel to avoid optical stream collisions from multiple sources. To check if the corresponding ejection channel is free, we use a separate optical control network (CN). Our CN is based on a Multiple-Writer Single-Reader optical bus [13].

The CN uses three kinds of control packets: request, acknowledgment (ACK) and negative acknowledgment (NACK). In the CN, each source node has a dedicated wavelength; a

node modulates a control packet on *its dedicated wavelength*. Therefore, a destination can accept multiple optical streams simultaneously as each will be on a different wavelength. Each source first sends a request packet over the CN waveguide connected to the desired destination. The source must wait for a response from the destination. An ACK is sent by the destination through the CN if the requested ejection channel is available (not being used by another source in the data network). Otherwise a NACK will be sent. Once the source receives the ACK, it begins sending data through the data network. If the destination is busy, a subsequent request packet will be sent after a back-off period which is determined by the average time required for sending a packet.

The CN is implemented as several waveguides (Figure 4). 16 nodes can accept data from each waveguide by utilizing splitters, which split an optical stream across different paths [14]. However, all nodes can send request packets on each waveguide. This structure helps reduce the waveguide crossings between laser power waveguide and CN waveguides. Hence, CN structure in an N -node QuT has $N/16$ waveguides and N wavelengths. By using more splitters on a waveguide, we can further reduce waveguide crossings. However, this increases the CN energy consumption, as a request packet is detected by more nodes per waveguide. As an example, if node 15 wants to send data to node $N-2$, it will modulate its control packet onto waveguide $N/16 - 1$. Each control packet has 6 bits: 4 bits for addressing 16 nodes and 2 bits for encoding the packet type. The receiver must have a $(N - 1) \times 6$ bit buffer to hold simultaneous requests from all other nodes.

III. METHODOLOGY

We compare a QuT with 64 and 128 nodes, in terms of latency and energy efficiency, against three alternative optical NoCs: λ -router [15], all-optical Spidergon [16] and Corona[17]. We use PhoenixSim, an event-driven simulator implementing optical NoCs [18]. It models the performance of electrical, optical and hybrid NoCs. PhoenixSim is based on OMNet++ [19], providing a component-based C++ library and discrete-event simulation framework.

λ -router [15] is an all-optical contention-free NoC. It uses passive microring resonators and routes optical streams in accordance with their wavelengths. We use an ideal λ -router which has infinite receiver-side buffering to eliminate the need for an arbitration mechanism. In practice, λ -router's performance will be limited when a more realistic buffer size is chosen. λ -router's scalability is limited by the number of wavelengths and switches. Since an N -node λ -router uses $N/2 \times (N - 1)$ switches, the number of switches increases quadratically with the number of nodes in network. λ -router with N nodes needs N distinct wavelength sets. A non-ideal λ -router would require a control network just as QuT does to avoid collisions at the destination due to limited receiver-side buffering. We choose to compare against an idealized version which allows us to assess the cost and performance degradation that occurs with the addition of a control network.

Optical Spidergon [16] uses passive optical components but has an electrical control network. We replace its electrical control network with the optical CN described in Section II-E to achieve better power and latency. In an all-optical Spidergon, each node has a dedicated but not unique wavelength, used by other nodes to modulate data. An all-optical Spidergon needs $N/2$ wavelength sets as each node is connected to its neighbors

Table I: Insertion Loss Parameters

Parameter	Value
C_{FtoW}	1 dB [21]
L_e	5 dB [21]
Waveguide propagation loss	1 dB/cm [22]
Passing through ring loss	0.5 dB [23]
Passing by ring loss	0.01 dB [23]
Waveguide bending loss	0.005 dB [23]
Waveguide crossing loss	0.12 dB [24]
Splitter loss	0.1 dB [14]
Receiver Sensitivity	-17 dBm [25]

in a ring and to the node that is $N/2$ nodes away. We choose Spidergon to compare to QuT since it is also a ring-based topology and uses cross links to reduce the NoC's diameter.

Corona [17] uses an optical crossbar with optical token-ring arbitration to permit a node to send data. It uses a Multiple-Writer-Single-Reader optical bus; when the number of nodes increases, both the waiting time for receiving a token and the network diameter increase. Therefore, we choose the best proposed token-ring arbitration scheme, slot-token-ring arbitration, as a control network of Corona [14]. The original Corona implementation uses a lot of wavelengths (64) to modulate data packets to improve the latency; however, this requires an impractically large number of microrings (one million) and high power consumption which rapidly increases with network size. To enable a consistent evaluation, we keep a constant flit-width of 1 byte across all networks. In Corona, the number of wavelengths equals the flit-width. Hence, we reduce the total number of wavelengths in Corona to eight distinct wavelengths. Therefore, our implementation improves Corona's power consumption and area by reducing the total number of microrings and waveguides, but its latency is increased. We choose Corona to compare with QuT since unlike QuT, its number of wavelengths is independent from the number of nodes, it does not need optical switches and it has a very simple architecture to implement the optical crossbar. We believe these three design points cover a range of interesting optical topologies that have been proposed in the literature.

We evaluate these NoCs with 64 and 128-nodes to compare their scalability. We assume the size of the die is the same for 64 and 128 nodes (225 mm^2). We keep the optical bandwidth constant for all-optical NoCs; we assign eight distinct wavelengths to each node. Therefore, each optical data stream is modulated on *eight* wavelengths assigned to the corresponding destination. However, each control packet is modulated on one wavelength assigned to the corresponding source. Since the control packets are small, the CN does not require large bandwidth. We assume a data packet size of 256 bits and 10Gb/s modulator and detector. We evaluate the topologies under the following synthetic traffic patterns: random, bitreverse, neighbor, tornado and hotspot in which a random node receives 30% of all requests and the remaining 70% is uniformly distributed. We evaluate these patterns with various offered loads. Offered load (α) is $T_{md}/(T_{md} + T_p)$, where T_{md} is the time that a packet takes to be transmitted through the network and T_p is the exponentially distributed inter-message gap [20].

IV. EVALUATION

A. Delay Evaluation

The total latency in an optical path includes: delay of the optical switches in the path, delay in bends and waveguide crossings, propagation delay in waveguides ($11.4 \text{ } \mu\text{s}/\mu\text{m}$),




Figure 5: Average Packet Latency for 64 nodes




Figure 6: Average Packet Latency for 128 nodes

modulator delay (23.8 ps) and detector delay (4.2 ps) [16]. We use aggressive delay parameters in 22nm. However, relative trends of our results will still hold with more conservative parameters. Also, improving modulator bit rate, e.g. from 10 to 40 Gb/s, can improve the latency, since modulating data to an optical stream is one of the major factors in latency. The total delay per packet includes the path delay through the control and data networks and the time the packet waits in the processor’s output buffer to be transferred into network.

Figures 5 and 6 show the average packet latency at $\alpha = 0.5$ for different traffic patterns for 64- and 128-node networks, respectively. Latency is reported in processor cycles. We assume a 5 GHz clock for the processors. QuT and Spidergon must check the status of the destination to avoid contention in the ejection channel. This destination checking and the ability to only accept one packet at a destination leads to an increase in latency for QuT and Spidergon, especially in high contention traffic such as hotspot. In hotspot traffic, 64- and 128-node QuTs saturate at $\alpha = 0.44$ and $\alpha = 0.38$, respectively because a destination can only receive from one source at a time.

The latency in 64- and 128-node QuT is up to 32% and 24% higher compared to 64- and 128-node Corona, respectively. Corona uses slot-token-ring arbitration to reserve a destination, which has better performance than QuT’s control network. However, it consumes more power and energy. The latency in a 128-node Corona is 11.5% higher than that of a 64-node Corona. In Corona, token waiting time and network diameter both increase as the number of nodes increases.

Waiting time in a processor’s output buffer coupled with the delay associated with modulating the packet, primarily contribute to the latency in QuT. When the network size grows from 64 to 128 nodes, the average optical path latency only increases by two cycles. Altogether, QuT’s latency does not rapidly increase for 128 nodes compared to 64 nodes, assuming the packet size, the number of wavelengths used to modulate a packet and the network offered load are the same.

λ -router has better latency compared to QuT since we assume that λ -router can accept data from all of the sources at a

time. Buffering for one packet per source is added to each destination. Hence, each node has 4064 bytes of buffering for 128-node λ -router and only 32 bytes for 128-node QuT, Spidergon and Corona. If the buffer size in λ -router is reduced, the network latency will increase sharply. The buffering overhead is not accounted in power and energy evaluations.

B. Power Evaluation

Total power consumption in an optical NoC includes off-chip laser power, on-chip microring heating power and electrical to optical (E/O) and optical to electrical (O/E) circuit power.

- **Laser Power:** Off-chip laser power is constant and independent of network traffic. This power is computed from the maximum optical insertion loss in the NoC. The insertion loss is the summation of: passing through the microring loss, passing by microring loss, waveguide crossing and bending loss and waveguide propagation loss. The laser power per λ , before entering to the chip, is calculated by $laserpower(\lambda) = -17 + IL + L_e + C_{FtoW}$, where -17 is the receiver sensitivity and L_e and C_{FtoW} represent laser efficiency and coupling coefficient, respectively. Finally, the total laser power is a multiplication of laser power per λ and total wavelengths in the NoC. To perform a fair comparison, we use the same optical parameters for all of the optical topologies, listed in Table I. Again, we consider aggressive parameters to show the full opportunity for power savings across these networks. More conservative parameters can increase the power consumption of optical NoCs, especially Corona.
- **Microring Heating Power:** We use thermally-tunable microrings. Extra power helps these microrings maintain their resonance wavelength when they experience temperature variation on chip. We use $20 \mu W/ring$, under typical conditions, when the rings in the system would experience a temperature range of 20K [26].
- **Total Power:** In QuT, Spidergon and Corona, the power of the CN must be factored in. Total power is the summation of total power consumption in the data and control networks.

Table II: Power dissipation for optical networks with 64 Nodes

	QuT & CN		Spidergon & CN		Corona & CN		λ -Router
	QuT	CN	Spidergon	CN	Corona	CN	
Max. Optical Loss (dB)	16.36	17	13.78	17	22.66	21.8	20.1814
Laser Power (λ) (mW)	3.44	4	1.9	4	14.66	12.02	8.282
Laser Power (mW)	440.32	256	486.4	256	117.28	769.28	4240.384
Ring Heating (mW)	901.12	85.76	1320.96	85.76	655.36	245.76	1955.84
Total Power (W)	1.69		2.15		1.8		6.196

Table III: Power dissipation for optical networks with 128 Nodes

	QuT & CN		Spidergon & CN		Corona & CN		λ -Router
	QuT	CN	Spidergon	CN	Corona	CN	
Max. Optical Loss (dB)	24.11	21	22.04	21	34.24	32.1	29.18
Laser Power (λ) (mW)	20.46	10	12.7	10	210.86	128.82	65.76
Laser Power (W)	5.24	1.28	6.5	1.28	1.69	16.49	67.34
Ring Heating (W)	3.44	0.346	5.26	0.346	2.62	0.983	7.844
Total Power (W)	10.31		13.39		21.78		75.18

Table IV: Characteristics of Optical Topologies

		QuT	Spidergon	CN(QuT& Spidergon)	λ -router	Corona	CN(Corona)
64-node	Num. of Wavelengths	128	256	64	512	8	64
	Num. of Microrings	45056	66048	4288	97792	32768	12288
128-node	Num. of Wavelengths	256	512	128	1024	8	128
	Num. of Microrings	172000	263000	17300	392192	131072	49152

The power dissipation for various optical networks with 64 and 128-nodes are shown in Tables II and III. Although QuT has higher optical loss compared with Spidergon, this is compensated by fewer wavelengths. QuT uses 256 wavelengths, while Spidergon needs 512 wavelengths for a 128-node network (Table IV). The required wavelengths are determined such that each topology has an 8-bit flit size. Microrings are sensitive to process and temperature variation. Therefore, they are a major source of faults in an optical NoC. An optical NoC should reduce the number of microrings to improve its reliability. QuT requires fewer microrings by 29.8%, 49.5%, 32.4% and 51.7% compared to a 64-node Spidergon and λ -router, and a 128-node Spidergon and λ -router, respectively (Table IV).

Power consumption increases for the 128-node optical NoCs due to increases in IL, the number of microrings and the total number of wavelengths. Reducing the number of wavelengths required to modulate data can reduce the power consumption. There is a trade-off between bandwidth and power consumption in an optical NoC. Future innovation in detector sensitivity will have a positive effect on the laser power [21]. For example by using a receiver that requires minimum power of -20 dBm, laser power of 128-node QuT is 2.63W, half of the laser power with a receiver requiring minimum power of -17 dBm (Table III). Among these optical NoCs, Corona's total number of wavelengths remains constant when the network size grows. However, Corona and its CN have the largest IL for 64 nodes among the other optical NoCs as shown in Table II. Therefore, increasing its IL, its CN IL and the number of wavelengths for 128 nodes, dramatically increases Corona's power consumption. The IL increases because of longer waveguide length and more waveguide crossings.

64-node QuT consumes up to 21%, 73% and 6% less power than Spidergon, λ -router and Corona, respectively. However, 128-node QuT achieves power reductions of up to 23%, 86.3% and 52.7% over Spidergon, λ -router and Corona, respectively.

C. Energy Evaluation

Energy consumption in an optical NoC includes the energy dissipation in the laser, micro-ring heating and back-end circuitry, E/O and O/E conversion. Based on the model in

Section IV-B, Figures 7a and 7b show the average energy-per-bit for different optical NoCs at $\alpha = 0.5$ for different traffic patterns. E/O and O/E consume 100 fJ/b [21].

In QuT and Spidergon, the average energy-per-bit includes the energy consumption in the data network and in the CN for multiple request packets per data packet if necessary. Energy consumption in Corona includes the energy consumption in both the data and control network. Since QuT has lower power dissipation and a smaller average optical path delay, it has lower energy-per-bit compared to Corona and λ -router. QuT achieves an energy-per-bit reduction of up to 69%, 92% over 128-node λ -router and Corona, respectively.

In Tornado traffic, a large number of packets are transferred between nodes with $N/2$ distance from each other; for a 128-node network, the average optical path delay increases. Thus, a data stream is sent through more optical switches and waveguides in QuT. Conversely, Spidergon benefits from this traffic pattern, since a data stream passes through less optical switches and waveguides; it has 34% reduction in energy-per-bit compared with QuT. However, QuT has fewer wavelengths and microrings and lower power consumption.

For hotspot traffic in QuT and Spidergon, a large number of requests are sent through the CN for a specific node. Due to destination contention, this results in a sharp increase in energy-per-bit (Figures 7a and 7b). At the saturation point in QuT and Spidergon, a small fraction of energy-per-bit is related to data network, e.g. the data network consumes 9.11% and 7.2% of the energy-per-bit in the 64-node QuT and Spidergon, respectively. Corona's CN always consumes energy even when the data network is idle, as Corona has a slot-tokening arbitration. Also, Corona's data network has an order of magnitude larger energy consumption compared with the CN. Thus, the average energy-per-bit of Corona at the saturation point in hotspot traffic does not increase significantly.

D. Throughput Evaluation

Effectively, QuT trades off throughput for lower cost and lower power consumption. Figure 8 shows the normalized average throughput-per-watt for optical NoCs. 64-node QuT has




Figure 7: Average Energy-per-bit. Traffic Patterns (excluding hotspot) use $\alpha = 0.5$.




Figure 8: Average Throughput-per-watt normalized to QuT. Traffic Patterns (excluding hotspot) use $\alpha = 0.5$

lower throughput-per-watt than Corona. However, since QuT consumes less power, 128-node QuT achieves throughput-per-watt improvements of up to 43%, 28% and 85% over 128-node Corona, Spidergon and λ -router, respectively, which indicates better scalability for QuT.

E. Area Evaluation

Area overhead in an optical NoC includes the area required for microrings, waveguides and detectors. Table V shows the normalized area for optical NoCs considered in this work. 128-node Spidergon and λ -router require 44% and 154% more optical area than QuT, since they have more microrings. λ -router uses more detectors to allow each node to accept data from different sources simultaneously. Although Corona has fewer microrings, it needs more waveguides and detectors compared to QuT. Therefore, its area is almost equal to QuT.

F. Discussion

QuT outperforms Spidergon, λ -router and Corona in terms of power dissipation and achieves better energy-per-bit consumption compared with λ -router and Corona. Although Spidergon slightly outperforms QuT in energy-per-bit, QuT has lower power consumption, fewer wavelengths (reduced by a factor of 2) and microrings. Due to lower resource requirements, QuT is easier to implement. QuT also uses fewer microrings and wavelengths compared with λ -router. To have a low power optical NoC, all of the optical factors mentioned in Section I should be considered. Having a topology only with smaller IL or fewer wavelengths does not lead to the best design choice in terms of power consumption.

Although QuT scales better compared to other optical NoCs, its power consumption does not scale well with node count. Therefore, without improving detector sensitivity or reducing optical bandwidth to compensate for the extra IL with increasing network size, QuT is not suitable for networks larger than 128 cores. However, if we use electrical or optical clustering, we can use QuT for systems larger than 128

Table V: Area overhead normalized to QuT

	QuT & CN	Spidergon & CN	λ -router	Corona & CN
64-node	1	1.38	2.42	1.06
128-node	1	1.44	2.54	1.03

cores. *Electrical clustering*: several cores can be clustered to connect to one optical node; e.g. a 128-node QuT with cluster size of 4 can be used to implement a 512-core system. *Optical clustering*: a single large optical NoC can be divided into several smaller parallel optical NoCs by using 3D stacking [21]; e.g. in 512-core system, instead of one 512-node QuT, we can use six 128-node QuTs to connect all of the cores. These optical NoCs are implemented in different optical layers to reduce the optical losses. Also, by using the electrical and optical clustering, we can connect 1024 cores with 6 parallel 128-node QuTs with electrical clustering size of 4.

V. RELATED WORK

Silicon photonics has generated significant interest among the research community as a solution to future interconnect obstacles. Several photonic NoCs [2][3][27][28][29][30] are implemented based on electro-optic broadband ring resonators. These optical NoCs use a circuit-switching approach and establish a path through the network before data is injected. A hybrid optical NoC combines an optical circuit-switched network for bulk message transmission with an electrical packet-switched NoC for control and short message exchange [2]. Each optical data transmission needs path setup and teardown, controlled by an electrical NoC. Thus, the performance of the electrical network and electro-optic broadband ring resonators and the use of electrical path-reservation affect the latency and power of the NoC. Large data messages can mitigate the cost of path reservation [3]; however, common packet sizes in many-core architectures are small (64-128 bytes to transmit a single cache line). Since QuT uses passive microring resonators, it has a simpler control network. The optical destination checking in QuT has significantly lower delay than electrical

path reservation resulting in superior performance compared to circuit-switching approaches.

Other optical networks [13][15][16][17][21][23][26] use passive microring resonators to route data based on wavelength. These optical NoCs eliminate the path reservation phase. Passive microrings consume less power than electro-optic broadband ring resonators. Firefly [13] is a hybrid hierarchical network architecture that employs a conventional electrical network for short/local data transmission and optical signaling for long/global communication. The crossbar-like optical NoC, is implemented through a Single-Writer-Multiple-Reader bus. The source node needs to broadcast a head flit to notify the destination node. R-3PO [21] utilizes an optical crossbar based on a Multi-Write-Single-Read optical bus. R-3PO uses token slot as a control network. Crossbar-like optical NoCs need more optical resources such as microrings, compared with QuT, when the network size or optical bandwidth is increased. As the network size increases, the power of the network increases rapidly. A photonic Clos [26] uses point-to-point channels to consume less power. The Clos uses more optical resources such as waveguides, compared to QuT. The number of waveguides in the Clos increases quadratically with the number of nodes which may limit its scalability. LumiNOC [31] is a lower power multi-stage design whose primary stage is optical. However, unlike QuT, it is not an all-optical NoC, since it uses electrical routers in its intermediate stage. At high injection rate, its performance is limited by electrical routers. Its arbitration policy is based on optical collision detection and dynamic channel scheduling techniques, which efficiently reuses optical resources for power efficiency.

VI. CONCLUSIONS

Integration of on-die photonics continues to pose challenges for both technology and architecture. NoC architects must identify clever solutions to overcome these technological limitations so that the promise of low-power, high-bandwidth on-die optical communication can be realized. To address challenges such as power, the number of microrings and wavelengths, we present QuT, an all-optical NoC using passive microrings. We propose a novel deterministic wavelength routing algorithm to optimize optical switches, reduce the number of wavelengths and microrings and lower the IL. QuT consumes less power and energy which allows it to scale better than state-of-art proposals. When the network size is increased, QuT is able to achieve lower power compared to Corona, λ -router and Spidergon.

ACKNOWLEDGEMENTS

The authors thank Joyce Poon, the members of the Enright Jerger research group and the anonymous reviewers for their valuable feedback and constructive suggestions to improve this work.

REFERENCES

- [1] J. Owens *et al.*, "Research challenges for on-chip interconnection networks," *Micro*, vol. 27, no. 5, pp. 96–108, 2007.
- [2] A. Shacham, K. Bergman, and L. Carloni, "Photonic networks-on-chip for future generations of chip multiprocessors," *IEEE Transactions on Computers*, vol. 57, no. 9, pp. 1246–1260, 2008.
- [3] J. Chan and K. Bergman, "Photonic interconnection network architectures using wavelength-selective spatial routing for chip-scale communications," *IEEE/OSA J. Optical Communications and Networking*, vol. 4, no. 3, pp. 189–201, 2012.

- [4] J. W. Goodman *et al.*, "Optical interconnections for VLSI systems," *Proceedings of the IEEE*, vol. 72, no. 7, pp. 850–866, 1984.
- [5] H.-S. Wang *et al.*, "Orion: a power-performance simulator for interconnection networks," in *MICRO*, 2002, pp. 294–305.
- [6] J. Kim *et al.*, "A low latency router supporting adaptivity for on-chip interconnects," in *DATE*, 2005, pp. 559–564.
- [7] N. Easley and L.-S. Peh, "High-level power analysis for on-chip networks," in *CASES*, 2004, pp. 104–115.
- [8] K. Preston *et al.*, "Performance guidelines for WDM interconnects based on silicon microring resonators," in *CLEO*, 2011, pp. 1–2.
- [9] A. Kovsh *et al.*, "Quantum dot comb-laser as efficient light source for silicon photonics," *Proc. SPIE*, vol. 6996, 2008.
- [10] Z. Lu *et al.*, "Low noise InAs/InP quantum dot c-band monolithic multiwavelength lasers for WDM-PONs," in *OFC/NFOEC*, 2009, p. JWA27.
- [11] P. Koonath, T. Indukuri, and B. Jalali, "Add-drop filters utilizing vertically coupled microdisk resonators in silicon," *Applied Physics Letters*, vol. 86, no. 9, 2005.
- [12] Q. Xu *et al.*, "Micrometre-scale silicon electro-optic modulator," *Nature*, vol. 435, no. 7040, pp. 325–327, 2005.
- [13] Y. Pan *et al.*, "Firefly: illuminating future network-on-chip with nanophotonics," in *ISCA*, 2009, pp. 429–440.
- [14] D. Vantrease, *Optical Tokens in Many-core Processors*. University of Wisconsin–Madison, 2010.
- [15] M. Briere *et al.*, "System level assessment of an optical NoC in an MPSoC platform," in *DATE*, 2007, pp. 1–6.
- [16] S. Koochi and S. Hessabi, "Scalable architecture for a contention-free optical network on-chip," *J. Parallel and Distributed Computing*, vol. 72, no. 11, 2012.
- [17] D. Vantrease *et al.*, "Corona: System implications of emerging nanophotonic technology," in *ISCA*, 2008, pp. 153–164.
- [18] J. Chan *et al.*, "PhoenixSim: a simulator for physical-layer analysis of chip-scale photonic interconnection networks," in *DATE*, 2010, pp. 691–696.
- [19] A. Varga, "The OMNeT++ discrete event simulation system," *Proc. of the European Simulation Multiconference*, June 2001.
- [20] A. Shacham, K. Bergman, and L. Carloni, "On the design of a photonic network-on-chip," in *NOCS*, 2007, pp. 53–64.
- [21] R. Morris, A. Kodi, and A. Louri, "Dynamic reconfiguration of 3D photonic networks-on-chip for maximizing performance and improving fault tolerance," in *MICRO*, 2012.
- [22] M. Lipson, "Guiding, modulating, and emitting light on silicon-challenges and opportunities," *J. Lightwave Technol.*, vol. 23, no. 12, p. 4222, Dec 2005.
- [23] M. J. Cianchetti, J. C. Kerekes, and D. H. Albonese, "Phastlane: a rapid transit optical routing network," *SIGARCH Comput. Archit. News*, vol. 37, no. 3, pp. 441–450, Jun. 2009.
- [24] W. Bogaerts *et al.*, "Silicon-on-insulator spectral filters fabricated with CMOS technology," *IEEE J. Selected Topics in Quantum Electronics*, vol. 16, no. 1, pp. 33–44, 2010.
- [25] X. Zheng *et al.*, "Ultra-efficient 10Gb/s hybrid integrated silicon photonic transmitter and receiver," *Opt. Express*, vol. 19, no. 6, pp. 5172–5186, Mar 2011.
- [26] A. Joshi *et al.*, "Silicon-photonic clos networks for global on-chip communication," in *NOCS*, 2009, pp. 124–133.
- [27] I. Artundo *et al.*, "Low-power reconfigurable network architecture for on-chip photonic interconnects," in *HOTI*, 2009, pp. 163–169.
- [28] J. Chan *et al.*, "Architectural exploration of chip-scale photonic interconnection network designs using physical-layer analysis," *J. Lightwave Technology*, vol. 28, no. 9, pp. 1305–1315, 2010.
- [29] H. Gu, J. Xu, and W. Zhang, "A low-power fat tree-based optical network-on-chip for multiprocessor system-on-chip," in *DATE*, 2009.
- [30] G. Hendry *et al.*, "Analysis of photonic networks for a chip multiprocessor using scientific applications," in *NOCS*, 2009.
- [31] C. Li *et al.*, "Luminoc: A power-efficient, high-performance, photonic network-on-chip for future parallel architectures," in *Proceedings of PACT*, 2012, pp. 421–422.