



Designing Low-power, Low-latency Networks-on-chip by Optimally Combining Electrical and Optical Links

DOI:
[10.1109/HPCA.2017.23](https://doi.org/10.1109/HPCA.2017.23)

Document Version
Accepted author manuscript

[Link to publication record in Manchester Research Explorer](#)

Citation for published version (APA):
Werner, S., Navaridas, J., & Luján, M. (2017). Designing Low-power, Low-latency Networks-on-chip by Optimally Combining Electrical and Optical Links. In *The 23rd IEEE Symposium on High Performance Computer Architecture (2017 IEEE International Symposium on High Performance Computer Architecture (HPCA))*. IEEE.
<https://doi.org/10.1109/HPCA.2017.23>

Published in:
The 23rd IEEE Symposium on High Performance Computer Architecture

Citing this paper
Please note that where the full-text provided on Manchester Research Explorer is the Author Accepted Manuscript or Proof version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version.

General rights
Copyright and moral rights for the publications made accessible in the Research Explorer are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Takedown policy
If you believe that this document breaches copyright please refer to the University of Manchester's Takedown Procedures [<http://man.ac.uk/04Y6Bo>] or contact openresearch@manchester.ac.uk providing relevant details, so we can investigate your claim.



Designing Low-power, Low-latency Networks-on-chip by Optimally Combining Electrical and Optical Links

Sebastian Werner, Javier Navaridas and Mikel Luján
The University of Manchester, Manchester, M13 9PL, UK
{sebastian.werner, javier.navaridas, mikel.lujan}@manchester.ac.uk

ABSTRACT

Optical on-chip communication is considered a promising candidate to overcome latency and energy bottlenecks of electrical interconnects. Although recently proposed hybrid Networks-on-chip (NoCs), which implement both electrical and optical links, improve power efficiency, they often fail to combine these two interconnect technologies efficiently and suffer from considerable laser power overheads caused by high-bandwidth optical links. We argue that these overheads can be avoided by inserting a higher quantity of low-bandwidth optical links in a topology, as this yields lower optical loss and in turn laser power. Moreover, when optimally combined with electrical links for short distances, this can be done without trading off latency.

We present the effectiveness of this concept with Lego, our hybrid, mesh-based NoC that provides high power efficiency by utilizing electrical links for local traffic, and low-bandwidth optical links for long distances. Electrical links are placed systematically to outweigh the serialization delay introduced by the optical links, simplify router microarchitecture, and allow to save optical resources. Our routing algorithm always chooses the link that offers the lowest latency and energy. Compared to state-of-the-art proposals, Lego increases throughput-per-watt by at least 40%, and lowers latency by 35% on average for synthetic traffic. On SPLASH-2/PARSEC workloads, Lego improves power efficiency by at least 37% (up to 3.5 \times).

1. INTRODUCTION

Many high-performance computing (HPC) systems are already equipped with chip multiprocessors that exhibit up to 100 cores [1, 2, 3] – a number expected to further increase. This shift rendered the on-chip network to be the limiting factor in terms of performance and power. Meanwhile, technology scaling has been creating energy and delay bottlenecks in electrical interconnect technologies [4]. It is commonly expected that electrical interconnects will not be able to satisfy the demands of future HPC applications [5].

Enabled by silicon photonics, optical on-chip communication has become a viable candidate to supplement or even replace electrical interconnects. Immense band-

width scalability through Dense Wavelength Division Multiplexing (DWDM), signal propagation of light, and almost distance-independent energy consumption are compelling properties to assume that optical interconnects might be a key disruptive technology for future many-core chips.

Optical Networks-on-chip (ONoCs) could either implement optical links only (all-optical) or combine them with electrical links (hybrid). In either case, their design is a challenging task which requires detailed knowledge of the properties of silicon photonic devices, power and throughput requirements of many-core chips, and a careful analysis of the relationship between optical bandwidth and power consumption. Nanophotonic devices and materials, as well as NoC architectures, are therefore hot research areas and essential to reach the full potential of ONoCs.

In this paper, we provide a detailed study of the trade-offs regarding latency and power consumption of optical and electrical links for current technologies, and identify the cases in which one interconnect technology should be preferred to another. Moreover, we conduct a systematic exploration of optical bandwidth vs. power consumption in DWDM links.

Based on our study, we propose Lego, a low-power hybrid ONoC design based on a design-friendly, mesh-based layout. We utilize electrical links for local communication with direct neighbors, where they provide low latency at low energy costs. Optical links are used in the mesh rows and columns for destinations residing at larger distances where they are superior to electrical links regarding energy and latency. We implement optical links with lower bandwidth than recent proposals and supplement them with electrical links to out-balance latency drawbacks introduced by serialization delays. Our routing algorithm minimizes the number of hops (at most two between any source-destination pair), and always chooses the path that provides the lowest energy and latency by performing distance-based routing to either route on electrical links, optical links, or a combination thereof. Minimizing link and router traversals, and always choosing the lowest energy and latency link, makes Lego highly efficient in terms of latency and dynamic power.




Figure 1: Single-Writer-Multiple-Reader Optical Bus

We make the following novel **contributions**:

- (i) We propose a novel hybrid NoC design Lego that saves power by combining electrical and optical links in an optimal, distance-based fashion without trading-off latency or throughput.
- (ii) Optical links are more abundant in Lego than in recent proposals but have lower bandwidth. This saves laser power without latency drawbacks when combined with electrical links for local traffic.
- (iii) Electrical links for neighbor traffic allow 1) energy-efficient short-distance communication, 2) a much simplified router microarchitecture and 3) savings of optical resources.
- (iv) Lego improves throughput-per-watt by up to 4x (40% on average) and packet latency by 35% on average (for synthetic traffic), and power efficiency by at least 37% (SPLASH-2/PARSEC).

2. OPTICAL VS. ELECTRICAL LINKS

Deciding on when to utilize optical and electrical links depends on a number of design aspects affecting latency and power consumption. In this section we discuss these implications for electrical and optical links, and identify their benefits and drawbacks.

Fig. 1 depicts a basic Single-Writer-Multiple-Reader optical bus (SWMR) in which Tile 0 sends and x number of tiles receive. A laser source provides wavelengths ($\lambda_1.. \lambda_n$) which are coupled into a waveguide. Tile 0 sends data by modulating on the n wavelengths, thus requiring a *modulator bank* with n modulators. To receive data, tiles implement *filter banks* – one ring filter for each λ . Optical data transmission includes optical data generation and serialization in the electrical-to-optical (E/O) backend circuitry, wavelength modulation, waveguide traversal, wavelength filtering, and optical-to-electrical data conversion (O/E) through detection and deserialization.

2.1 Power Consumption

2.1.1 Optical Links

Static power is the main contributor to the total power consumption in optical interconnects and consists of laser and ring heater power. The latter is required to mitigate temperature variations and post-manufacturing geometric mismatches of microring resonators. Ring heater power used to be prohibitively high in NoCs on a larger scale (> 10000 rings) [6]; however, significant research efforts in ring tuning techniques in recent years lead to a decrease in ring heating down to $20\mu\text{W}/\text{ring}$

or even $5\mu\text{W}/\text{ring}$ for some technologies [7] [8]. Athermal microring resonators are also a hot research topic which would reduce temperature dependency to a tolerable level that would cancel the need for ring heating of microrings [9] [10]. Ring heating power is therefore likely to be manageable in the near future.

Laser power depends on losses of silicon photonic devices, which require the laser source to provide more output power to drive all receivers at satisfactory bit-error rates. The optical path that introduces the highest insertion loss (IL_{max}) determines the output power per wavelength. Current laser efficiencies and device losses require a very careful design of optical links to avoid excessive laser power [11]. Although devices are evolving, there is no clear roadmap for nanophotonics. The NoC layout should therefore aim to minimize IL_{max} .

Besides IL_{max} , laser power depends on the number of wavelengths provided by the laser source and the number of readers/detectors it has to drive, depicted in Fig. 2 and Fig. 3, respectively. We modeled laser power with DSENT [12] using IL parameters listed in Tab. 2. Fig. 2 demonstrates the relation between laser power and number of wavelengths in an 10mm Single-Writer-Single-Reader (SWSR) bus. Laser power grows exponentially with increasing number of wavelengths since more wavelengths do not only have a direct effect on the laser source itself, but also increase IL_{max} as they lead to more microrings that are passed by wavelengths ('ring-through' loss), as well as crosstalk noise which significantly increases with the number of wavelengths [11]. From a power perspective, it is therefore more efficient, for instance, to implement two SWSR buses with 16λ each rather than one 32λ bus. A larger number of low-bandwidth links in a topology could thus be more power-efficient than a small number of high-bandwidth links shared by a number of nodes.

These effects are further aggravated when there is more than one reader attached to an optical bus, as shown in Fig. 3, as even more microrings lead to even higher IL_{max} . Increasing numbers of readers lead to higher laser power since the laser source has to drive more photodetectors, and is increasingly critical with higher bandwidth. It is therefore desirable to keep the number of readers and wavelengths of an optical bus as low as possible; however, as these two aspects also determine bandwidth/throughput, it is crucial consider these aspects carefully for low-power, low-latency designs.

2.1.2 Electrical Links

Global electrical wires have become increasingly energy-hungry in many-core architectures [4], as they require repeaters, regenerators or buffers to provide satisfactory signal integrity and latency, with increasing energy consumption for longer link lengths. Fig. 4 shows the difference in energy per 64-bit flit over an electrical and optical link with increasing link length, modeled in DSENT with a 22nm technology. For short distances, the electrical link is more energy-efficient as it does not require E/O and O/E conversion circuitry. However, for link lengths $> 1\text{mm}$, the almost distance-independent




Figure 2: Laser Power vs. Num. of Wavelengths in a SWSR bus




Figure 3: Laser Power vs. Num. of Receivers in a SWMR bus

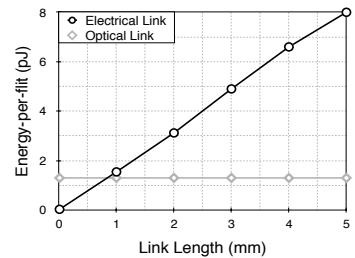


Figure 4: Energy for transmitting a 64-bit flit

Table 1: Latency of a 64-bit flit on an optical link. For simplicity: $t_{prop} = 1$

Number of Wavelengths	Serialization Degree	Delay (E/O + t_{prop} + O/E)
32λ	1:1	3 (1 + 1 + 1)
16λ	2:1	4 (2 + 1 + 1)
8λ	4:1	6 (4 + 1 + 1)
4λ	8:1	10 (8 + 1 + 1)

energy consumption of optical data transmission outperforms electrical links. From an energy perspective, it is therefore only beneficial to utilize electrical links for destinations $< \sim 1\text{mm}$. For instance, in a 64-core chip, tile widths/lengths are often between 1-2mm for common die sizes of 225mm^2 . This would mean that only communication to direct neighbors should be electrical. Besides, router traversal of a 64-bit flit in 22nm at 5Ghz requires 2pJ, which is similar to the energy needed to traverse a link of 1.3mm - further emphasizing the impact electrical links have on total energy consumption.

2.2 Latency

Electrical signal propagation takes 131ps/mm in an optimally repeated wire at 22nm [13]. At 5Ghz, one hop over an electrical link in a NoC is therefore commonly accepted to take one clock cycle (we note that this is also subject to clock frequency, layout, final link lengths, etc.). Optical links, on the other hand, require E/O and O/E conversions and on-the-fly propagation delay (t_{prop}), which take at least one clock cycle each (3 cycles in total). However, optical links leverage the signal propagation of light (11.4ps/mm for current technologies [14]), which is particularly beneficial for long-distance communication, especially because optical links, as opposed to electrical, do not require pipelining and do not introduce further distance-related latencies. For instance, with 11.4ps/mm propagation delay, data can be sent within one clock cycle to any core located at distances $< 17.5\text{mm}$ assuming a core clock frequency of 5Ghz (200ps clock cycle).

Although all optical components add to the optical delay (modulator (3.1ps), detector(0.22ps), E/O(9.5ps) and O/E(4.0ps) [15]) the major contributor is data modulation, i.e. the time it takes to serialize a packet based on the available bandwidth. This is outlined in Tab. 1, which lists the impact on the delay of a 64-bit flit with different number of wavelengths, assuming link propagation delay of 1 cycle for simplicity, 10Gb/s modu-

lators and 5Ghz clock frequency. These values are an important guideline to trade-off power and latency. For instance, increasing link bandwidth from 16λ to 32λ decreases latency only by one clock cycle, but more than 8x introduces too much latency for too little power benefits.

To result in minimum packet latencies, these delays must be compared to electrical delay. Although electrical links do not need E/O and O/E conversions, the only energy-efficient way of reaching distant cores is through several hops in a topology, which includes router delay. Router traversal delay depends on the clock frequency, where high clock frequencies of 5Ghz may need up to 5 pipeline stages, as in Intel's TeraFLOPS design [2]. If we assume aggressively pipelined routers that can be traversed in two clock cycles (assuming enough link bandwidth), one hop would take 3 cycles. While this delay adds up for each additional hop to reach a destination, hardly any delay is added on optical links when the distance increases (assuming direct connections). From this perspective, optical links are superior when destinations are at 2-hop distances or further away in a topology when optical bandwidth is at least 8λ .

3. NETWORK ARCHITECTURE

This section introduces Lego, in which we apply the findings of the previous section to minimize laser power, dynamic power, and packet latency, by

- Minimizing the number of wavelengths available for optical data transmission while still staying superior to electrical links in terms of latency.
- Minimizing the number of readers in optical links.
- Minimizing energy consumption by 1) using electrical links only for 1-hop distances, 2) using optical links otherwise, and 3) keeping the total number of link traversals low by allowing paths of at most of 2 hops for any source-destination pair.

3.1 Topology

Lego arranges nodes in a 2D mesh as this allows an efficient layout of the optical links, as we will discuss in Section 3.1.2. Each node is connected to its direct neighbors via electrical links, illustrated in green in Fig. 5/6. Moreover, they are connected to every

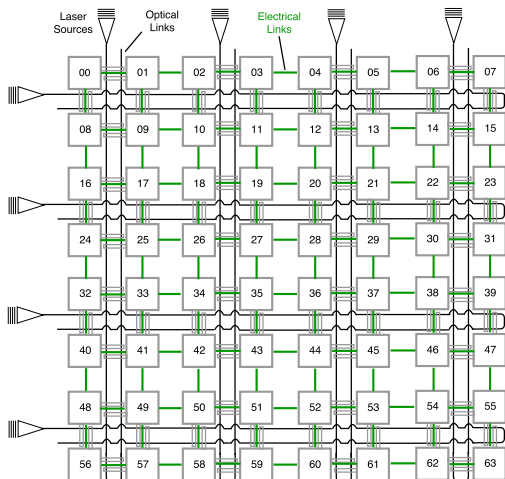


Figure 5: Lego8: 8 router groups (four row/column) connect 16 nodes each

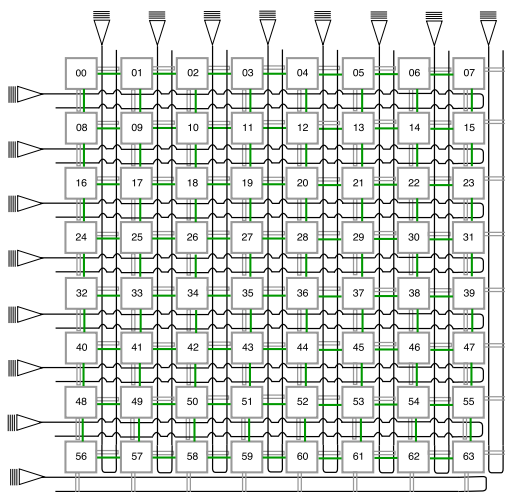


Figure 6: Lego16: 16 router groups (eight row/column) connect 8 nodes each

node in their **optical router group** with an optical link. We study two different variants of Lego, Lego8 and Lego16, which both utilize the same routing algorithm and router architecture, but provide a different number and arrangement of optical links. This allows us to study the trade-off of varying levels of bisection bandwidth and power consumption. Lego8 has four optical router groups in its rows and columns (8 in total), with 16 nodes belonging to one router group. In Lego16, each row/column is an optical router group (16 in total), with 8 nodes belonging to one router group. The optical layout is similar to LumiNoC [16], however, LumiNoC deploys optical links only, resulting in a different routing algorithm and router microarchitecture.

3.1.1 Optical Router Group

All nodes belonging to the same optical router group are connected to each other in a crossbar fashion, with a distinct SWMR bus for each node. Fig. 7 gives a close-up to an optical router group with 16 nodes, as implemented in Lego8. For simplicity, the figure only shows

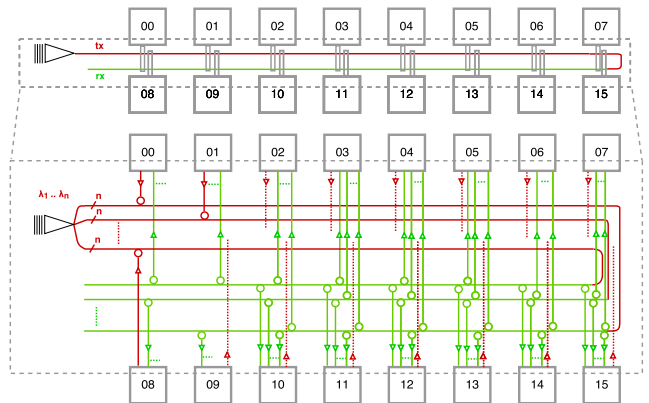


Figure 7: Optical router group layout (here for Lego8). Control network has the same layout but is omitted for simplicity. Nodes modulate data on the tx path (red) and receive on the rx path (green).

the SWMR buses of a few nodes. Every other node owns an equal, separate bus for sending. We chose a U-shaped layout (like in [16]) of the optical links which allows nodes to reach every other node by modulating data on the transmit side of the link (red). All receiving nodes attached to the bus filter out the optical data on the receive side (green). Optical router groups with 8 nodes like in Lego16 have the same layout like in Fig. 7, just without the lower row of nodes (08-15).

In our topology, nodes do not need modulators, ring filters and detectors for communication with its direct neighbors because they are connected to them electrically. This reduces the total number of microrings, and thus area and ring heater power. For instance, in Fig. 7, node 01 and 08 have no ring filters on the SWMR bus of node 00, because our routing algorithm will choose electrical links for neighbor traffic. The same applies to every other node in our topology and its neighbors. The difference between Lego8 and Lego16 topology-wise is that the former provides fewer optical link groups, but in each group twice as many nodes are connected. This allows to reach more nodes in one-hop distance, decreasing zero-load latency. However, this provides less bandwidth as each node has to use the same optical link to reach a larger number of nodes. Lego16 has twice as many optical link groups, but connects half the number of nodes, which decreases the number of nodes in one-hop distance, and thus zero-load latency, but also increases bandwidth which might be beneficial for certain workloads and injection rates. More optical links, however, also lead to higher laser and ring heater power. In Section 4, we study these different topological considerations and their effect on latency and power.

Control Network.

Based on our study in Section 2, we aim to keep the number of receivers on the SWMR buses as low as possible to result in low laser power. So far, for N number of nodes attached to each router group and n number of the sender's direct neighbors, the number of readers the coupled laser source of each SWMR bus has to drive is

$N - n - 1$ (-1 for the sender itself, N being either 8 or 16). This would result in high laser power overheads as Fig. 3 illustrates. We therefore implement a parallel, low-overhead control network for each SWMR bus, similar to the 'reservation-assisted' SWMR bus in [17]. The purpose of this control network is to control the receiver's ring filters so that, at any time, only one node on the receive side of the bus filters out modulated data. This significantly reduces laser power as it sets the number of readers the laser source has to drive to one. This is enabled by ring heaters capable of tuning microrings by shifting their resonance wavelength(s) so that they respond to particular wavelengths, or detune them to let wavelengths pass without filtering them - allowing to switch on and off microrings as desired. Tuning speeds have been subject to extensive research [18, 19, 20]. Recent studies found microrings to optically stabilize in less than 100ps [18], and total tuning times of at most 500ps [19] - 2-3 clock cycles at 5Ghz clock frequency. As E/O conversion and optical signal propagation takes at least one cycle, we pipeline data transmission by letting the sender start sending one cycle before ring tuning is finished. This leads to one clock cycle of ring tuning/detuning delay, with a reasonable assumption of 400ps ring tuning latency.

Each node in the router group owns one SWMR control bus for its SWMR data bus to realize this functionality. Both buses have the same layout, merely the number of wavelengths, and thus modulators and ring filters, differ. Transmitting data over optical links thus obeys the following process:

1. Initially all nodes are detuned and do not filter the wavelengths.
2. When a node wants to send data, it first sends out a control packet containing the destination node and packet size on the control network to all receivers.
3. The destination node tunes in and all other nodes keep their filters on the data bus detuned. The packet size indicates the duration of which ring filters have to stay tuned/detuned.
4. The sender starts transmitting its data.

Control packets are very small and therefore only require low bandwidth. We support two packet sizes, 64 bit and 576 bit, as they are present in common CPU architectures for miss request/coherence traffic and cache line transfers, respectively. Depending on the router group size, either 3 or 4 bits are needed to encode the destination ID (< 8 (Lego16) or < 16 (Lego8), sender and direct neighbors are not part of the possible destinations). Adding one bit to encode the packet size, this makes 4 bits / 5 bits, which can be modulated in one clock cycle with 2/3 wavelengths, respectively. Assuming one clock cycle for packet processing and ring tuning each, this would result in a latency overhead of 5 clock cycles (on-the-fly delay of optical signals in Lego is always 1 clock cycle as the distance between two nodes in

a router group is always $< 17.5mm$). As we will show later, the latency and power overhead introduced by the control network is negligible compared to the laser power savings it provides. We note, however, that this latency overhead slightly distorts our latency analysis in Section 2, and would make electrical links the faster medium for 2-hop distances, too, and not just for direct neighbor traffic. We still stick to our restriction to use electrical links only for one hop traffic as our analysis in DSENT showed that 2-hops in the electrical domain require 3.5x the energy of 1-hop in the optical domain. We leave the study of whether this large overhead is worth the latency benefits in this case to future work.

3.1.2 Layout

Implementing ONoCs with silicon photonics requires a careful consideration of the implications of layout and device technologies. We target low laser power by providing a larger number of low-bandwidth links rather than few high-bandwidth links. It is important to note that this approach decreases laser power only if the higher number of waveguides does not lead to a higher number of waveguide crossings, which could increase IL_{max} and possibly diminish some of the power savings. In our layout, waveguide crossings occur when optical links located in the columns cross the ones in the rows. For that reason, we assume 3D-integration with the optical circuitry of row and column router groups placed separate photonic layers to eliminate in-plane waveguide crossings [21].

Depending on the topology, an optical router group has either 8 (Lego16) or 16 (Lego8) nodes attached to it, which requires 16 or 32 waveguides for the data and control network. Although current technologies allow waveguide dimensions of 520nm width [11], sufficient clearance is needed to avoid optical signal interference. As microring resonators have a diameter of $5\mu m$ [22], we assume a waveguide pitch of $15\mu m$, leaving $5\mu m$ clearance [16]. For 16 and 32 waveguides per router group, this layout requires $< 0.25mm$ and $< 0.5mm$, respectively, for the optical links in the rows and column. With a die size of $225mm^2$, this would allow the conventional tile sizes of $1mm^2$ while providing sufficient area for interfacing and placement of the photonic devices in the topology's rows and columns.

Our mesh-based layout is not only benign to VLSI floor-planning, but also omits the need for a large number of laser sources and allows to place laser sources on the edges of the chip. This is important as chip packaging is one of the major cost factors of silicon photonic chips due to expensive coupling of the off-chip laser source. Therefore, designs are likely to have to oblige to tight packaging constraints which have to be taken into account by designers.

3.2 Routing Algorithm

Our routing algorithm aims to minimize link traversals and always chooses the link that offers the lowest energy and latency to keep dynamic power and latency

as low as possible. Therefore, based on the relative position of the sender and destination, the former either sends on an optical link, electrical link, or a combination thereof. We classify routing into four cases, demonstrated in Fig. 8 - 11 in a 6x6 Lego8 design for simplicity (green links indicate hops over electrical links and blue links over optical):

1. **Case 1** (Fig. 8): Source and destination node are direct neighbors in the 2D mesh. In this case, the source sends its packet directly to the destination using electrical mesh links.
2. **Case 2** (Fig. 9): Source and destination node are in the same column or row group, but not direct neighbors. In this case, the source node uses its column/row optical link to send data directly to the destination node (e.g. node 00 to 07, and node 14 to 23).
3. **Case 3**: Source and destination node are neither in the same row nor column group. In this case, the source will first use its optical row link to send to the node that resides in the same column group as the destination node. Once the packet is received by the intermediate node, two possibilities exist:
 - (a) **Case 3.1** (Fig. 10): The destination node is a direct neighbor. The intermediate node proceeds by sending data to the destination via the electrical mesh link (e.g. node 00 to 26, and node 14 to 29).
 - (b) **Case 3.2** (Fig. 11): The destination node is not a direct neighbor. The intermediate node then proceeds by sending data to the destination via its optical column link (e.g. node 00 to 35).

Each node merely needs to compute its own and its destination’s position, and possibly one intermediate node in case the destination is not directly reachable. To provide our architectural simplifications to the electrical network, these are the only existing routing scenarios (e.g. it is not possible to use first an electrical and then an optical link). In the worst case (Case 3.2), optical data transmission is performed twice, and router traversal three times. This leads to high efficiency both in terms of packet latency and dynamic power consumption. Moreover, the electrical mesh network is simplified substantially because no actual routing needs to be performed on the electrical links: A node uses a mesh link *only* if the packet’s destination is on the other side of this link, making routing computation dispensable since each incoming packet on this link is forwarded to the local tile.

3.3 Router Microarchitecture

In Lego, each router has to handle data communication through both its electrical and optical ports. We illustrate our router microarchitecture proposal in Fig. 12. Depending on the relative position of the router in the

mesh (either at the border or center), the number of output ports for the electrical mesh links varies. In both Lego8 and Lego16, each node has two output ports for the optical links in its row and column group.

Optical input and output ports implement the E/O and O/E signal conversion circuitry prior to the input buffers. Based on our routing algorithm, incoming data on the mesh links is always intended for the local core. Therefore, no routing computation has to be performed. It suffices to store incoming flits in buffers and multiplex them to the output port leading to the local core. This requires a small crossbar for arbitration between packets incoming from the mesh links, and those that were received on the other ports. However, this greatly simplifies the router’s crossbar from a 7x7 to 3x7 crossbar, leading to a much lower power and area footprint than conventional mesh routers. For every other input port, routing computation and switch allocation is executed in the conventional way.

4. EVALUATION

4.1 Experimental Set-up

We compare Lego to a wide range of the most recent NoC proposals. Several ONoCs utilizing optical links for long-distance and electrical links for short-distance communication have been proposed in recent literature. We chose to compare Lego to the most competitive designs that come closest to our goal of combining electrical and optical links in the most efficient manner, i.e. Atac [23], Firefly [17], and Meteor [15]. In addition, we compare Lego to the state-of-the-art, low-power all-optical NoC QuT [7], and LumiNoC [16]. To outline the benefits of implementing optical links, we also add an electrical baseline 2D Mesh.

We use DSENT [12] for area and power estimations, and Sniper [24] for performance and energy modeling of SPLASH-2 [25] and PARSEC [26] applications with the *sim-large* input set. Results are measured during the parallel phase of the applications after caches have been warmed up. Sniper is configured according to Xeon X550 Gainestown chip multiprocessor [27], and uses private L1I and L1D caches (32kB) and a shared LLC (16MB), with memory controllers placed on the top and bottom rows. We use a 22nm technology, 5Ghz clock, 10Gb/s modulators/detectors, and a die size of $225mm^2$ with square tiles. For synthetic traffic, we use the cycle-accurate simulator HNOCS [28], and assume a data packet size of 256 bits and flit size of 64 bit. We study our NoC topologies for 64 nodes and assume an 8x8 layout.

Atac [23] consists of a 2D electrical mesh that is overlaid by an optical network (ONET). In the 64-node version, each node is connected to the ONET, which is a bundle of 32 SWMR links that carry 64 wavelengths each. Packets for destinations less than four hops away are sent on the electrical mesh, and on the ONET otherwise. Both electrical and optical links are 64-bit wide. Atac+ [10] improves Atac by adding a more performance-efficient star network and an adap-

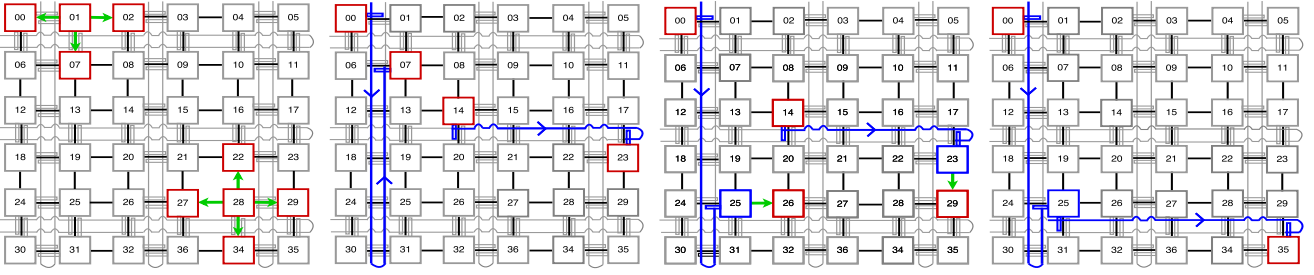


Figure 8: Routing Case 1 Figure 9: Routing Case 2 Figure 10: Routing Case 3.1 Figure 11: Routing Case 3.2

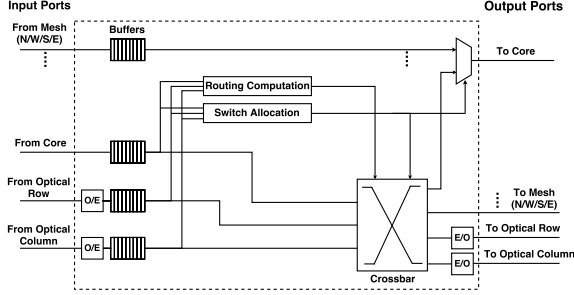


Figure 12: Router Architecture

tive Ge laser that allows to adjust the output power to the traffic demands. As the latter is a technological advancement, and we are focusing on architectural improvements in this paper, we only adapt the star network to the traditional Atac design.

Firefly [17] divides the 64 nodes into four similar sized clusters of 16 nodes each. Within each cluster, four hub routers form an electrical 2D mesh, with four nodes concentrated at each hub. Each of these four hub routers has a dual in every other cluster, with which they are connected optically. Hubs use the electrical mesh to send to destinations residing in the same cluster, and optical links to their duals for inter-cluster communication. Optical links are implemented as **reservation-assisted SWMR** buses (R-SWMR), as introduced earlier. Therefore, prior to optical data transmission, control packets are exchanged and rings tuned/detuned. We assume enough optical link bandwidth so that control packets (4 bits) and data packets are modulated in one clock cycle, i.e. 2λ on the control channels, and 32λ on the data channels. Electrical links are 64-bit wide.

Meteor [15], similar to Atac, implements a 2D electrical mesh and overlays it with an optical network. However, in Meteor, there are only four optical hubs through which the ONET can be accessed. Photonic Regions of Influence (PRI) determine the grouping of nodes to the hubs. With an 8×8 layout, their study shows that grouping 16 nodes to each PRI is the most efficient design variant. We divide the 8×8 layout into four square 4×4 submeshes and place the hub router in the middle of each submesh for the highest efficiency. If the destination node is closer than the node’s PRI hub, it will send the data packet over the electrical mesh. Otherwise, it will route the packet to the hub, which will then send the packet optically to the PRI region of the destination node, which will then forward the packet to the

destination using the mesh. The ONET connects the PRI hubs using four Multiple-Writer-Multiple-Reader (MWMR) buses, on which each hub can send on a 128-bit wide link. We assume 64-bit wide electrical links.

QuT [7] is a low-power, all-optical NoC that uses passive microring resonators to route optical signals according to their wavelength. Senders modulate data on the wavelength set that is assigned to the destination they want to address. For N nodes, QuT uses $N/4$ wavelength sets for addressing to reduce the number of microrings and laser power. As every destination has one ejection channel, a separate control network is required to resolve contention at the destinations, implemented by exchanging control messages prior to data transmission. The control network is implemented with MWSR buses. Control packets are modulated on one wavelength, and data packets on 8 wavelengths [7].

LumiNoC [16] has the same topology as Lego16, but does not implement electrical links and uses optical data transmission only. If the destination is in the same row or column, it can be reached in one hop. Otherwise, XY optical routing is performed like in routing case 3.2 (Fig. 11). LumiNoC implements a MWMR bus combined with an arbitration mechanism to share optical bandwidth. In this paper, we are interested in the power and performance benefits of topologies that optimally combine electrical and optical links. Therefore, we implement the router groups in LumiNoC like in Lego, with 8λ on each link, without the arbitration functionality. This allows us to study how much benefits we can get by combining electrical and optical links, rather than using optical links only. We note that this arbitration mechanism can be adopted to Lego as well. **EMesh** is a conventional 2D electrical mesh network with 64-bit wide links and 5Ghz clock. Routers implement XY-routing and use wormhole switching. We assume an optimistic design with aggressively pipelined routers and three cycles per hop: two within each router and one for traversing a link. We chose the electrical mesh, as it is the *de facto* standard in industry [29] and constitutes a baseline electrical NoC.

We study **Lego8** and **Lego16** at two different bandwidth levels, i.e. optical links carrying 8λ (**Lego8_8 λ** , **Lego16_8 λ**), and 16λ (**Lego8_16 λ** , **Lego16_16 λ**). Increasing the bandwidth to 16λ halves serialization delay without excessive power overheads (see Tab. 1 and Fig. 2). Control links carry 2λ (**Lego16**) and 3λ (**Lego8**). We assume 64-bit wide electrical links.

4.2 Packet Latency

4.2.1 Synthetic Workloads

Fig. 13 illustrates the average packet latency for some of the traditional synthetic traffic patterns. We report latency in processor cycles. In Hotspot traffic, 80% of the nodes send all of their traffic randomly to 20% of the nodes, while the rest distributes their traffic uniformly. In Neighbor traffic, each node sends packets randomly to one of its neighbors.

Compared to Atac, Firefly, and Meteor, which utilize electrical links for local traffic and optical links for distant traffic, both Lego topologies tremendously decrease packet latency across all traffic patterns. Only in neighbor traffic, Meteor shows similar latencies. At a fairly low injection rate of 1Tbps, both Lego topologies manage to decrease packet latency by 50% on most patterns. At the same time, our topologies at least double throughput on all patterns.

Compared to LumiNoC, the electrical links inserted in our topology for local traffic, along with our novel routing algorithm, show to be a considerable improvement. Packet latency is decreased by at least 30% on all patterns, and substantial throughput gains can be observed. Even Lego8, which provides fewer optical links than LumiNoC and Lego16, demonstrates these improvements for most patterns (apart from Bit-Reversed and Bit-Complement).

QuT shows very constant latency and throughput across all patterns as the distance to the destination does not have a large impact on packet latency for optical signal propagation, as opposed to congestion and contention resolution, which are the main contributor to latency in QuT. Therefore, QuT performs particularly poorly in Hotspot traffic. On all traffic patterns, both Lego topologies decrease latency (up to 50%, 20% on average at 1Tbps). Moreover, they provide fair throughput gains (apart from the Bit-Permutation patterns, where QuT provides slightly more throughput).

For most traffic patterns, Lego8 and Lego16 show similar throughput and latency values, with Lego16 having slightly lower latencies for Uniform Random, Bit-Reserved, and Bit-Complement. Only for Bit-Reversed and Bit-Complement, Lego16 offers twice the throughput of Lego8, while having similar throughput levels for every other traffic pattern. It is interesting to observe that there is no workload on which Lego seems to perform particularly poorly compared to the alternative NoCs. Moreover, both Lego designs are very efficient for Neighbor traffic, which is a desired property for NoCs as it supports the current shift to data centric computing in large-scale many-core systems, where software tends to exploit spatial locality through near-data processing [30]. Increasing the bandwidth from 8λ to 16λ provides only slight throughput and latency gains for both Lego8 and Lego16. Section 4.3 evaluates whether these gains justify the entailed power overheads.

4.2.2 PARSEC/SPLASH-2 Workloads

Fig. 14 shows the latency results of a range of SPLASH-

2 and PARSEC applications normalized to Lego8_8λ. Apart from Atac, every alternative NoC has a higher packet latency than both Lego implementations. Atac has optical links between all source-destination pairs which provides high bandwidth, but also high power consumption as we will see in the following section. Compared to Lego8_8λ, Firefly and Meteor exhibit a slight latency overhead of 14% and 6%, respectively. The same applies to QuT and EMesh, which have a latency overhead of 56% and 30%, respectively. Lego saves up to 40% by inserting electrical links compared to LumiNoC.

In addition, Lego8_8λ has the highest latency of all Lego networks. Similar to synthetic traffic, Lego16_8λ lowers packet latency compared to Lego8_8λ. On average, 13% fewer cycles are required. Increasing the optical link bandwidth in both Lego topologies leads to larger savings in packet latency than for synthetic traffic, with 20% lower packet latency for both Lego8_16λ and Lego16_16λ.

4.3 Power Consumption

Fig. 16 depicts the power breakdown normalized to Lego8_8λ. The power values are the average power consumption across all synthetic traffic patterns before network saturation at 1Tbps. We report ring heater power for thermally-tunable microring resonators that require 20μW/ring for a typical on-chip temperature range of 20K [22]. Encasing the photonic die in a thermal insulator can further decrease this value to ~5μW [8]. In this work, we use 20μW/ring to have a more pessimistic assessment of the silicon photonic technology.

Compared to the other locally-electrical, globally-optical NoCs Atac, Firefly, and Meteor, both Lego implementations decrease power consumption significantly. Even when increasing the bandwidth to 16λ, Lego exhibits lower power consumption than the majority of alternative NoCs. Atac provides optical links for each node, which is very inefficient in terms of laser and ring heater power (4x the power consumption of Lego8_8λ). Meteor only has very few optical links in its topology and thus low laser and ring heater power. However, it heavily relies on the underlying electrical mesh for most source-destination pairs, resulting in higher dynamic power. Our results show a 49% and 46% power overhead compared to Lego8_8λ and Lego16_8λ, respectively. Firefly implements fewer optical links than Lego, but with larger optical bandwidth, leading to higher static optical power. Its topology also leads to more electrical link and router traversals than Lego, which is reflected in higher electrical dynamic power. In total, Firefly increases power consumption by 62% and 59% compared to Lego8_8λ and Lego16_8λ, respectively.

Being all-optical, QuT exhibits according static optical power overheads. As only optical communication is performed, dynamic power is very low due to the highly energy-efficient optical communication. Nevertheless, with power consumption overheads of 76% and 73% compared to Lego8_8λ and Lego16_8λ, its power requirements are high.

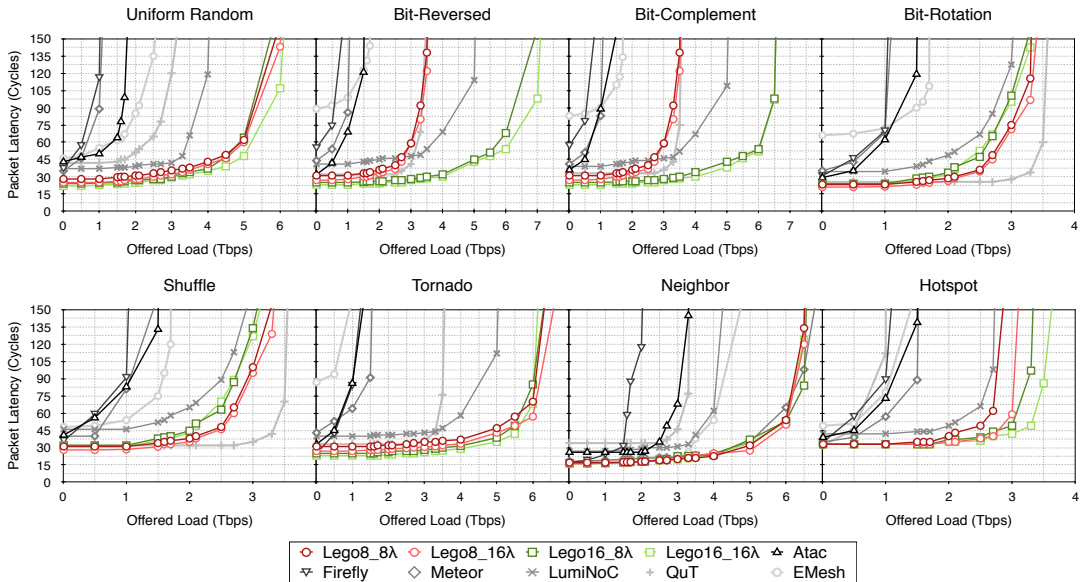


Figure 13: Average packet latency on synthetic workloads




Figure 14: Average packet latency on SPLASH-2/PARSEC benchmarks normalized to *Lego8_8λ*

Compared to LumiNoC, electrical links in Lego increase energy efficiency as they allow considering locality and provide fewer total link traversals. Given that we do not have to perform routing computation when a packet is received over an electrical link, Lego saves dynamic router power. Moreover, Lego has fewer optical links and microrings because they are not required for communicating with direct neighbors. These architectural aspects lead to a 22% and 19% power overhead of LumiNoC compared to *Lego8_8λ* and *Lego16_8λ*, respectively. Compared to the electrical mesh, these savings are 74% and 71%.

Increasing the optical bandwidth to 16λ increases the power consumption by $\sim 50\%$ for both Lego topologies, which shows the susceptibility of static optical power to optical bandwidth. However, slight power savings can still be observed.

4.3.1 Throughput-per-Watt

We show the throughput-per-watt (TPW) of all NoCs for all synthetic traffic patterns in Fig. 15. Both Lego designs dramatically increase TPW compared to each of the alternative NoCs for each traffic pattern. LumiNoC is the closest competitor, and only provides 71% of the TPW of *Lego8_8λ*. For every other NoC, *Lego8_8λ* at least doubles the TPW, proving its high energy ef-

Table 2: Insertion Loss Parameters

Parameter	Value
Laser efficiency	0.25 [31]
Coupler	1 dB [31]
Ring: Through	0.01 dB [12]
Ring: Drop	1 dB [12]
Waveguide Bending	0.005 dB [7]
Waveguide propagation	0.1 dB/mm [32]
Waveguide crossing	0.05 dB [22]
Splitter	0.1dB [7]
Photodetector loss	1 dB [12]

iciency. *Lego16_8λ* shows the highest TPW and increases TPW compared to *Lego8_8λ* by $\sim 8\%$. Increasing the optical bandwidth of Lego from 8λ to 16λ still provides good TPW improvements, but does not seem to offer sufficient throughput to justify its power overheads. In both Lego topologies, doubling the link bandwidth leads to $\sim 25\%$ lower TPW.

4.3.2 Power-Delay-Product

We calculate the power-delay-product (PDP) for the SPLASH-2/PARSEC applications by multiplying the average packet latency with the consumed power. We present the results in Fig. 17. Our Lego topologies




Figure 15: Normalized Throughput-per-Watt

show significant improvements in PDP compared to the alternative NoCs. The large power overhead of Atac makes it the least energy-efficient of all designs. QuT is less energy-efficient than the remaining NoCs as it has large static, data-independent optical power consumption and many applications of the SPLASH-2 and PARSEC benchmark suite have fairly low injection rates over the whole duration of executing time [33]. Meteor is the closest competitor and has a 46% higher PDP than *Lego8_8λ*. The difference in PDP of *Lego8_8λ* and *Lego16_8λ* are marginal. Increasing the bandwidth to 16λ raises PDP by 25% in both cases, but they are both still considerably more efficient than the other NoCs.

4.4 Area

The area breakdowns are shown in Fig. 18. Lego basically trades off latency and power for area compared to most NoCs. Given their large abundance of optical components, Atac and QuT require more area than *Lego8_8λ* and *Lego16_8λ*. Meteor, LumiNoC, Firefly, and EMesh save 22%, 28%, 38%, and 30% in area compared to *Lego8_8λ*, respectively, which is not negligible; however, both *Lego8_8λ* and *Lego16_8λ* outperform these alternative NoCs in most of the other metrics. It is commonly assumed that area constraints are going to be less significant in future on-chip network designs than power, especially with shrinking transistors sizes and 3D-Integration enabling to integrate silicon photonic devices on a separate layer. Nevertheless, the area overheads of Lego are not of prohibitive extent and we believe are more than justified considering the latency, throughput, and power gains provided by Lego.

5. RELATED WORK

As one of the most promising emerging technologies to overcome the energy and performance limitations of metal wires, optical NoCs have gained large interest in the research community, ranging from improving photonic device technologies [8, 34, 35], thermal management [36, 37], and adaptive laser sources [38, 39] to novel NoC architectures that makes use of this nascent technology in the most efficient manner. The paper at hand contributes to this field by proposing a novel, low-power ONoC architecture that studies the trade-offs of electrical and optical links more carefully and utilizes them more efficiently.

A number of NoC designs have been proposed which we classify into 1) all-optical NoCs that use optical data transmission only, and 2) hybrid NoCs that combine

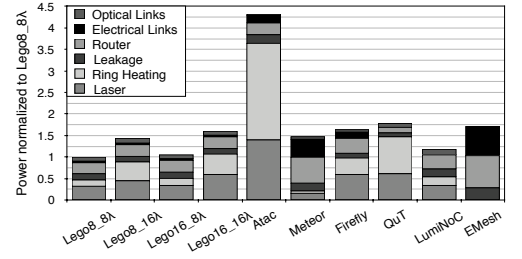


Figure 16: Power breakdown

electrical and optical links.

Passive microring resonator based ONoCs are referred to as Wavelength-Routed ONoCs (WRONoCs). Non-blocking WRONoC topologies [40, 41, 42, 43] provide simultaneous switching capability from each sender to each receiver; however, this requires $(N - 1)$ filters at each node, leading to very limited scalability in terms of power as the number of rings is proportional to ring heater power. To tackle this, several WRONoCs [14, 7, 44] have been proposed that require only one filter-detector pair at each node and resolve contention by using a control network like QuT. This decreases power as fewer microrings are used, but fairly large static optical power is still present in these NoCs, especially for higher core counts.

Hybrid NoCs are more likely to be adopted in near future than all-optical NoCs because of the currently large static power overheads of optical interconnects. A large number of interesting hybrid optical NoCs has been proposed in recent literature [45, 46, 47, 17, 22, 48, 49, 50, 51, 52, 31, 10, 16, 15]. Phastlane [47] combines a packet-switched mesh network with an optical, contention-free crossbar to transmit cache lines over several hops in one cycle. The silicon-photonic Clos network (PClos) [22] uses point-to-point optical channels for low-energy, long distance data transmission. Only the router in the intermediate clos stages and the links from/to the cores to the output/input routers are electrical, optical links are used in all other stages. It could thus also be considered an all-optical NoC. BLOCON [51] is a bufferless implementation of PClos that features a scheduling algorithm and path allocation scheme for managing routing in the Clos. It provides low latency and high throughput, but also has higher ring heater and laser power compared to PClos. FlexiShare [48] deploys a channel sharing architecture to reduce channel over-provisioning and thus laser power at the cost of less throughput and additional arbitration channels. PROPEL [49] combines an optical crossbar with an electrical mesh. The number of wavelengths required in the NoC equals the number of nodes in the topology. R-3PO [31] is a 3D NoC utilizing an optical crossbar with a token-based control network to handle accesses. Just like WRONoCs, optical crossbars have limited scalability as the number of microrings, and thus ring heater power, increases rapidly for larger network sizes. Also, high-bandwidth optical links are shared by a number of nodes, which leads to higher IL and laser power. Lego improves these designs by providing smaller crossbars with lower throughput channels,

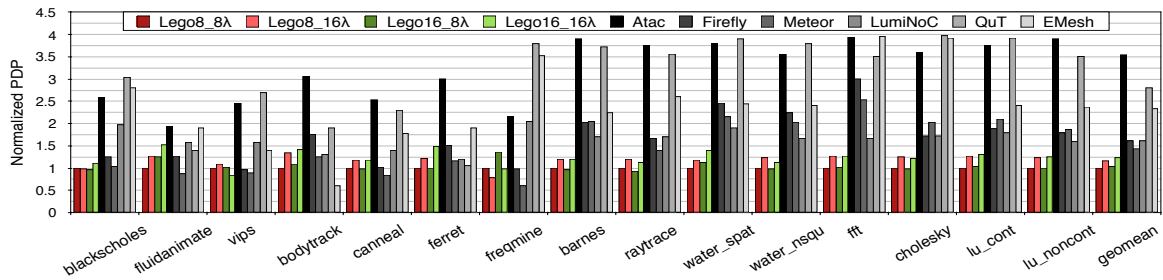


Figure 17: Power-Delay-Product normalized to *Lego8_8λ*




Figure 18: Normalized Area Breakdown

and supplements them more efficiently with electrical links to outbalance throughput drawbacks. Utilizing 3D stacking technology for the sake of mitigating losses caused by waveguide crossings has also been successfully studied [50, 53], which is why we adapted this approach in Lego, too. Iris [52] combines optical links, a dielectric antenna-array-based broadcast network, and a circuit-switched electrical mesh network. Lego does not impose any limitations on implementing emerging technologies, such as adaptive laser sources. Channel allocation schemes can also be implemented in Lego, which would further lower laser power [16].

6. CONCLUSION

We present Lego, a hybrid ONoC topology that decreases power consumption without latency and throughput drawbacks by efficiently combining electrical and optical links. Our novel routing algorithm utilizes electrical and optical links based on the distance to the destination, which always picks the technology that provides the lowest energy and latency to transmit a packet. Low-bandwidth optical links offer satisfactory throughput when deployed for sufficiently large distances. Our evaluation results prove the effectiveness of Lego by exhibiting large savings in throughput-per-watt and power-delay-product compared to various alternative NoC proposals. We intend to extend our study in the future to larger network sizes to study scalability. Moreover, the effect of applying channel allocation schemes and adaptive laser sources in Lego to further decrease laser power are also interesting aspects to investigate.

7. ACKNOWLEDGEMENTS

This research was conducted with support from the UK Engineering and Physical Sciences Research Council (EPSRC) PAMELA EP/K008730/1 and the Euro-

pean Union’s Horizon 2020 research and innovation programme under grant agreement No 671553 (ExaNeSt). Dr. Luján is supported by a Royal Society University Research Fellowship.

8. REFERENCES

- [1] G. Chrysos, “Intel® xeon phi™ coprocessor-the architecture,” *Intel Whitepaper*, 2014.
- [2] S. Vangal, J. Howard, G. Ruhl, S. Dighe, H. Wilson, J. Tschanz, D. Finan, P. Iyer, A. Singh, T. Jacob, S. Jain, S. Venkataraman, Y. Hoskote, and N. Bordkar, “An 80-tile 1.28 tflops network-on-chip in 65nm cmos,” in *ISSCC*, pp. 98–99, IEEE, 2007.
- [3] T. Corporation, “Tilera multicore processors.” <http://www.tilera.com/products/processors>, 2007.
- [4] M. A. Anders, “High-performance energy-efficient noc fabrics: Evolution and future challenges,” in *NOCS*, pp. i–i, IEEE, 2014.
- [5] I. O’Connor and G. Nicolescu, *Integrated optical interconnect architectures for embedded systems*. Springer Science & Business Media, 2012.
- [6] M. Georgas, J. Leu, B. Moss, C. Sun, and V. Stojanović, “Addressing link-level design tradeoffs for integrated photonic interconnects,” in *CICC*, pp. 1–8, IEEE, 2011.
- [7] P. K. Hamedani, N. E. Jerger, and S. Hessabi, “QuT: A low-power optical network-on-chip,” in *NOCS*, pp. 80–87, IEEE, 2014.
- [8] Y. Demir and N. Hardavellas, “Parka: Thermally insulated nanophotonic interconnects,” in *NOCS*, p. 1, ACM, 2015.
- [9] K. Padmaraju and K. Bergman, “Resolving the thermal challenges for silicon microring resonator devices,” *Nanophotonics*, vol. 3, pp. 269–281, 2014.
- [10] G. Kurian, C. Sun, C.-H. O. Chen, J. E. Miller, J. Michel, L. Wei, D. A. Antoniadis, L.-S. Peh, L. Kimerling, V. Stojanovic, and A. Agarwal, “Cross-layer energy and performance evaluation of a nanophotonic manycore processor system using real application workloads,” in *IPDPS*, pp. 1117–1130, IEEE, 2012.
- [11] K. Bergman, L. P. Carloni, A. Biberman, J. Chan, and G. Hendry, *Photonic network-on-chip design*. Springer, 2014.
- [12] C. Sun, C.-H. O. Chen, G. Kurian, L. Wei, J. Miller, A. Agarwal, L.-S. Peh, and V. Stojanovic, “Dsnet-a tool connecting emerging photonics with electronics for opto-electronic networks-on-chip modeling,” in *NOCS*, pp. 201–210, IEEE, 2012.
- [13] R. Ho, “Wire scaling and trends,” a presentation at mto darpa meeting,” *Sun Microsystems Laboratories, Jackson Hole, WY*, 2006.
- [14] S. Koohi and S. Hessabi, “Scalable architecture for a contention-free optical network on-chip,” *Journal of Parallel and Distributed Computing*, vol. 72, pp. 1493–1506, 2012.
- [15] S. Bahirat and S. Pasricha, “Meteor: hybrid photonic ring-mesh network-on-chip for multicore architectures,” *TECS*, vol. 13, p. 116, 2014.

- [16] C. Li, M. Browning, P. V. Gratz, and S. Palermo, "Luminoc: A power-efficient, high-performance, photonic network-on-chip," *TCAD-CEDA*, vol. 33, no. 6, pp. 826–838, 2014.
- [17] Y. Pan, P. Kumar, J. Kim, G. Memik, Y. Zhang, and A. Choudhary, "Firefly: illuminating future network-on-chip with nanophotonics," in *ACM SIGARCH Computer Architecture News*, vol. 37, pp. 429–440, ACM, 2009.
- [18] A. Shacham, K. Bergman, and L. P. Carloni, "Photonic networks-on-chip for future generations of chip multiprocessors," *Transactions on Computers*, vol. 57, pp. 1246–1260, 2008.
- [19] E. Peter, A. Thomas, A. Dhawan, and S. R. Sarangi, "Active microring based tunable optical power splitters," *Optics Communications*, vol. 359, pp. 311–315, 2016.
- [20] J. Tang, M. Li, Y. Yang, S. Sun, N. Shi, W. Li, and N. Zhu, "High-speed tunable broadband microwave photonics phase shifter based on an active microring resonator," in *WOCC*, pp. 1–3, IEEE, 2016.
- [21] 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," *JETC*, vol. 7, p. 7, 2011.
- [22] A. Joshi, C. Batten, Y.-J. Kwon, S. Beamer, I. Shamim, K. Asanovic, and V. Stojanovic, "Silicon-photonics networks for global on-chip communication," in *NOCS*, pp. 124–133, IEEE Computer Society, 2009.
- [23] G. Kurian, J. E. Miller, J. Psota, J. Eastep, J. Liu, J. Michel, L. C. Kimerling, and A. Agarwal, "Atac: a 1000-core cache-coherent processor with on-chip optical network," in *PACT*, pp. 477–488, ACM, 2010.
- [24] W. Heirman, T. Carlson, and L. Eeckhout, "Sniper: scalable and accurate parallel multi-core simulation," in *ACACES*, pp. 91–94, HiPEAC, 2012.
- [25] S. C. Woo, M. Ohara, E. Torrie, J. P. Singh, and A. Gupta, "The splash-2 programs: Characterization and methodological considerations," in *ACM SIGARCH Computer Architecture News*, pp. 24–36, ACM, 1995.
- [26] C. Bienia, S. Kumar, J. P. Singh, and K. Li, "The parsec benchmark suite: characterization and architectural implications," in *PACT*, pp. 72–81, ACM, 2008.
- [27] Intel, "Intel Xeon Processor 5500 Series." <http://ark.intel.com/products/37106>.
- [28] Y. Ben-Itzhak, E. Zahavi, I. Cidon, and A. Kolodny, "Hnocs: modular open-source simulator for heterogeneous nocs," in *SAMOS*, pp. 51–57, IEEE, 2012.
- [29] C. Ramey, "Tile-gx100 manycore processor: Acceleration interfaces and architecture," in *HC*, 2011.
- [30] R. Balasubramonian, J. Chang, T. Manning, J. H. Moreno, R. Murphy, R. Nair, and S. Swanson, "Near-data processing: Insights from a micro-46 workshop," *IEEE Micro*, vol. 34, pp. 36–42, 2014.
- [31] R. Morris, A. K. Kodi, and A. Louri, "Dynamic reconfiguration of 3d photonic networks-on-chip for maximizing performance and improving fault tolerance," in *MICRO*, pp. 282–293, IEEE Computer Society, 2012.
- [32] M. Lipson, "Guiding, modulating, and emitting light on silicon-challenges and opportunities," *J. Lightw. Technol.*, vol. 23, p. 4222, 2005.
- [33] J. Lee, C. Nicopoulos, S. J. Park, M. Swaminathan, and J. Kim, "Do we need wide flits in networks-on-chip?," in *ISVLSI*, pp. 2–7, IEEE, 2013.
- [34] G. T. Reed, G. Z. Mashanovich, F. Y. Gardes, M. Nedeljkovic, Y. Hu, D. J. Thomson, K. Li, P. R. Wilson, S.-W. Chen, and S. S. Hsu, "Recent breakthroughs in carrier depletion based silicon optical modulators," *Nanophotonics*, vol. 3, pp. 229–245, 2014.
- [35] H. Subbaraman, X. Xu, A. Hosseini, X. Zhang, Y. Zhang, D. Kwong, and R. T. Chen, "Recent advances in silicon-based passive and active optical interconnects," *Optics express*, vol. 23, pp. 2487–2511, 2015.
- [36] T. Zhang, J. L. Abellán, A. Joshi, and A. K. Coskun, "Thermal management of manycore systems with silicon-photonics networks," in *DATE*, pp. 1–6, IEEE, 2014.
- [37] H. Li, A. Fourmigue, S. Le Beux, X. Letartre, I. O'Connor, and G. Nicolescu, "Thermal aware design method for vcsel-based on-chip optical interconnect," in *DATE*, pp. 1120–1125, EDA Consortium, 2015.
- [38] C. Chen and A. Joshi, "Runtime management of laser power in silicon-photonics multibus noc architecture," *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 19, pp. 3700713–3700713, 2013.
- [39] Y. Demir and N. Hardavellas, "Ecolaser: an adaptive laser control for energy-efficient on-chip photonic interconnects," in *ISLPED*, pp. 3–8, ACM, 2014.
- [40] M. Briere, B. Girodias, Y. Bouchebaba, G. Nicolescu, F. Mileyville, F. Gaffiot, and I. O'Connor, "System level assessment of an optical noc in an mp soc platform," in *DATE*, pp. 1084–1089, EDA Consortium, 2007.
- [41] D. Vantrease, R. Schreiber, M. Monchiero, M. McLaren, N. P. Jouppi, M. Fiorentino, A. Davis, N. Binkert, R. G. Beausoleil, and J. H. Ahn, "Corona: System implications of emerging nanophotonic technology," in *ACM SIGARCH Computer Architecture News*, vol. 36, pp. 153–164, IEEE Computer Society, 2008.
- [42] S. Le Beux, J. Trajkovic, I. O'Connor, G. Nicolescu, G. Bois, and P. Paulin, "Optical ring network-on-chip (ornoc): Architecture and design methodology," in *DATE*, pp. 1–6, IEEE, 2011.
- [43] L. Ramini, P. Grani, S. Bartolini, and D. Bertozzi, "Contrasting wavelength-routed optical noc topologies for power-efficient 3d-stacked multicore processors using physical-layer analysis," in *DATE*, pp. 1589–1594, EDA Consortium, 2013.
- [44] S. Werner, J. Navaridas, and M. Luján, "Amon: An advanced mesh-like optical noc," in *HOTI*, pp. 52–59, IEEE, 2015.
- [45] N. Kirman, M. Kirman, R. K. Dokania, J. F. Martinez, A. B. Apsel, M. A. Watkins, and D. H. Albonese, "Leveraging optical technology in future bus-based chip multiprocessors," in *MICRO*, pp. 492–503, IEEE Computer Society, 2006.
- [46] S. Pasricha and N. Dutt, "Orb: An on-chip optical ring bus communication architecture for multi-processor systems-on-chip," in *ASP-DAC*, pp. 789–794, IEEE Computer Society Press, 2008.
- [47] M. J. Cianchetti, J. C. Kerekes, and D. H. Albonese, "Phastlane: a rapid transit optical routing network," *ACM SIGARCH Computer Architecture News*, vol. 37, pp. 441–450, 2009.
- [48] Y. Pan, J. Kim, and G. Memik, "Flexishare: Channel sharing for an energy-efficient nanophotonic crossbar," in *HPCA*, pp. 1–12, IEEE, 2010.
- [49] R. Morris and A. K. Kodi, "Exploring the design of 64-and 256-core power efficient nanophotonic interconnect," *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 16, pp. 1386–1393, 2010.
- [50] X. Zhang and A. Louri, "A multilayer nanophotonic interconnection network for on-chip many-core communications," in *DAC*, pp. 156–161, ACM, 2010.
- [51] Y.-H. Kao and H. J. Chao, "Blocon: A bufferless photonic clos network-on-chip architecture," in *NOCS*, pp. 81–88, IEEE, 2011.
- [52] Z. Li, M. Mohamed, X. Chen, H. Zhou, A. Mickelson, L. Shang, and M. Vachharajani, "Iris: A hybrid nanophotonic network design for high-performance and low-power on-chip communication," *JETC*, p. 8, 2011.
- [53] S. Pasricha and S. Bahirat, "Opal: A multi-layer hybrid photonic noc for 3d ics," in *ASP-DAC*, pp. 345–350, 2011.