



Efficient Sharing of Optical Resources in Low-Power Optical Networks-on-Chip

DOI:
[10.1364/JOCN.9.000364](https://doi.org/10.1364/JOCN.9.000364)

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). Efficient Sharing of Optical Resources in Low-Power Optical Networks-on-Chip. *Journal of Optical Communications and Networking*, 9(5), 364-374.
<https://doi.org/10.1364/JOCN.9.000364>

Published in:
Journal of Optical Communications and Networking

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 uml.scholarlycommunications@manchester.ac.uk providing relevant details, so we can investigate your claim.



Efficient Sharing of Optical Resources in Low-Power Optical Networks-on-Chip

Sebastian Werner, Javier Navaridas and Mikel Luján,
School of Computer Science, The University of Manchester, Manchester, UK
Email: {sebastian.werner, javier.navaridas, mikel.lujan}@manchester.ac.uk

Abstract—With the ever-growing core counts in modern computing systems, NoCs consume an increasing part of the power budget due to bandwidth and power density limitations of electrical interconnects. To maintain performance and power scaling, alternative technologies are required, with silicon photonics being a promising candidate thanks to high-bandwidth, low-energy data transmission. In order to get the best of silicon photonics, sophisticated network designs are required to minimize static power overheads. In this paper, we propose Amon, a low-power ONoC that decreases number of μ Rings, wavelengths and path losses to reduce power consumption. Amon performs destination checking prior to data transmission on an underlying control network, allowing the sharing of optical bandwidth. Compared to a wide range of state-of-the-art optical, hybrid, and electrical NoCs, Amon improves Throughput-per-Watt by at least 23% (up to 70%), while reducing power without latency overheads on both synthetic and realistic applications. For aggressive optical technology parameters, Amon considerably outperforms all alternative NoCs in terms of power, outlining its increasing superiority as technology matures.

Index Terms—Optical Network-on-Chip, Silicon Photonics, On-chip Interconnects



1 INTRODUCTION

As the trend towards many-core chips with increasing core counts continues [1]–[3], Networks-on-Chip (NoCs) interconnecting on-chip components have become an essential part of chip design. However, technological limitations of electrical interconnects cause energy and delay bottlenecks that prevent the performance and power scaling of future many-core devices. Recent breakthroughs in silicon photonics have enabled optical on-chip communication, which provides high-bandwidth, low-latency data transmission and distance-independent energy consumption - intriguing attributes for many-core chips with increasing core counts and throughput demands. While already an established technology in conventional computer networks, optical data transmission on-chip is a nascent research field with a number of design obstacles still to be overcome through novel network architectures.

The major obstacles of optical NoCs (ONoCs) for widespread commercial adoption are the static power overheads of nanophotonic components, in particular laser and ring heater power. Laser power is the output power required at the light source to ensure low error rates. It depends on the number of wavelengths and readers, and on the highest insertion loss path (IL_{max}). Ring heating is required to mitigate the effects of temperature variations and device mismatches on microring resonators (μ Rings), the basic building blocks for modulators and filters. Designers have to be aware of these implications to minimize their power overheads.

We previously proposed Amon, our low-power ONoC [4] which routes packets based on their mod-

ulated wavelength. Wavelengths are assigned to each destination for addressing and are shared between nodes to reduce their number, and so laser power. Amon's topology provides collision-free paths to avoid data corruption of two nodes sending simultaneously on the same wavelengths. Compared to previous ONoCs, Amon decreases laser power by minimizing IL_{max} on the paths, and μ Ring heater power by reducing the number of μ Rings.

In this paper, we further develop Amon by proposing an efficient optical layout, providing a thorough performance and power evaluation, and studying the impact of advances in nanophotonic technologies. We make the following novel **contributions**: (i) Discuss the implications of layout and placement of on-chip components, and propose an *optimized layout* for Amon which improves upon our previous design. (ii) Enhance the control network for destination checking. (iii) Analyze the potential of *aggressive* silicon photonic technology parameters. (iv) Compare Amon to *state-of-the-art* optical, hybrid and electrical NoCs and evaluate benefits and trade-offs regarding power consumption, area, and performance. (v) Found that Amon improves *Throughput-per-Watt* by up to 70% (at least 23%), while providing slight latency reductions on both synthetic and PARSEC/SPLASH-2 applications.

2 RELATED WORK

The task of efficiently implementing optical links in NoCs has gained large interest in the research community. Approaches can broadly be classified into designs based on 1) electro-optic broadband μ Ring resonators, 2) hybrid NoCs implementing both electrical and optical

links, and 3) Wavelength-Routed ONoCs (WRONoCs) based on passive μ Rings.

Electro-optic broadband μ Ring designs [5]–[8] use a circuit-switching approach in which broadband rings are tuned to according wavelengths to establish a path prior to data transmission. Latency introduced by path reservation and release can be mitigated by using large messages. However, these are uncommon in NoCs (e.g. a cache line is typically 64–128 bytes).

Hybrid optical NoCs aim to combine electrical and optical links in the most efficient manner. A large number of interesting hybrid optical NoCs has recently been proposed [9]–[22]. The most recent and competitive hybrid NoC proposals—Atac [23], Firefly [12], and Meteor [22]—are compared to our design proposal and will be discussed in more detail in section 4.1.

WRONoCs based on passive μ Rings route optical streams through the network based on their wavelengths. Non-blocking WRONoC topologies [24]–[27] provide simultaneous switching capability from each sender to each receiver. However, this requires $(N - 1)$ filter-detector pairs (for N nodes) at each node, leading to a very limited scalability: μ Ring heater power and area increase quadratically with N . To tackle this issue, several blocking WRONoCs [28], [29]—including Amon [4]—have been proposed. These require only one filter-detector pair at each node and resolve contention by using a control network (CN). This decreases μ Ring power and area. Furthermore, laser power is reduced by sharing wavelengths for addressing and providing collision-free paths in the network. CoNoC [28] first introduced this principle and has $N/2$ wavelengths for addressing. QuT [29] significantly improves upon CoNoC by splitting the wavelengths to $N/4$. Our proposal, Amon, further improves upon QuT by providing a more efficient layout that needs fewer μ Rings for routing, and has shorter paths for less insertion loss.

3 AMON: NETWORK ARCHITECTURE

In Amon, each node has one ejection channel for receiving data. This tremendously decreases the number of receivers, but also means that only one packet can be received at a time which requires an arbitration mechanism for access management. This is handled by a separate CN on which nodes have to request to destinations prior to data transmission on the data NoC.

3.1 Data Network

3.1.1 Logical Topology

In Amon, we split the number of wavelength sets (λ) for addressing into $N/4$. The wavelength set size determines the number of wavelengths on which senders can modulate their data (e.g. 8λ , 16λ , etc.), and in turn network throughput. Fig. 1 shows the topology of a data network with 48 nodes. The plot shows node id (unique) and wavelength assignment for addressing (unique within each submesh but reused across them). All links are bidirectional (i.e., require two waveguides)

unless indicated by arrows. As four nodes share the same λ -set as address, up to four nodes could be sending data on the same λ -set at a time. Collision-free paths are always guaranteed in Amon as there is only one λ -set address per submesh, and physically separated waveguides to send data to each submesh. Amon Relies on the following kinds of links:

Mesh links (shown in black): are the regular links within a 2D mesh topology.

Intermesh links (shown in blue): are used for data transmission to nodes in other submeshes. They use separate waveguides parallel to the Mesh links.

Amon's data network can be scaled very flexibly as it does not require a square mesh. Its only limitation is that all submeshes must have the same size due to row and column connections through Intermesh links. Therefore, the 48-node Amon in Fig. 1 could also be composed by 3×4 , 6×2 or 2×6 submeshes, if needed.

3.1.2 WDM Routing

Amon uses passive μ Rings to route optical streams based on their wavelength. Each μ Ring responds to one particular wavelength, determined by its geometry and dimensioning. Therefore, all routing paths in Amon are *predefined* and effectively perform static routing. From a sender's perspective, a destination can either be

In the same submesh: packets will be injected into the local submesh and routed within using static dimension-ordered routing (DOR).

In a different submesh: packets will be injected onto the Intermesh links leading to the destination's submesh. Once in the desired submesh, the local submesh is used as above.

This simple behavior is directly implemented in hardware (see the router architecture below) and requires very simple computation at injection time: $\lambda\text{-set} = \text{destination} \bmod N/4$ and $\text{submesh} = \text{destination} \text{ divided by } N/4$. With typical power of 2 number of nodes, this can be further simplified to bit operations (the 2 most significant bits define the submesh and the remaining bits define the wavelength). Fig. 3 shows routing examples on a 64-node Amon:

(i) **Red path**—Node 43 sends to Node 10: N_{43} modulates data on its Intermesh link to the left and on the λ -set λ_{10} to address N_{10} . Once the optical stream reaches N_{10} 's submesh, the μ Rings at N_6 will route the stream down in Y-direction to N_{10} . Node 10 has μ Ring filters that respond to λ_{10} at its ejection channel to successfully receive data.

(ii) **Green path**—Node 1 sends to Node 64: N_1 modules data on λ_{16} to address N_{64} on the Intermesh link leading to N_{64} 's submesh, i.e. the Intermesh link to the right. Again, once the submesh is reached, the μ Ring filters performing DOR lead the optical stream to N_{64} , where it will finally be ejected.

As we will show later in more detail when we discuss the physical layout, switching optical streams with μ Ring filters significantly adds to IL_{max} , and thus laser

power. Therefore, one goal is to minimize the amount of switching required between any source-destination pair. We achieve this by implementing DOR as follows: In (i), the Intermesh link leading from the lower-left to the upper-left submesh does not stop at n_5 : It goes straight through the submesh, all the way to N_8 . This way, there only needs to be μ Rings switching the signal either up or down, depending on the λ -set. E.g. at N_6 , μ Ring filters are implemented to route λ_{10} down to N_{10} . Similarly, in (ii), the Intermesh link leading from the upper-left to lower-right submesh goes all the way through to N_{52} in the latter submesh. At N_{52} , the according μ Rings routes λ_{16} down towards N_{64} . Similar μ Rings implemented in N_{49} , N_{50} and N_{51} route streams down to the their respective columns. For clarity, this means that Intermesh links from a submesh are extended to be the submesh links in another one (see Sec. 3.1.4 for more details). The logical distinction is made for the sake of simplicity.

3.1.3 Router Microarchitecture

Routers at each node have to provide μ Ring filters to implement the switching performed in the routing algorithm by dropping according wavelengths from one waveguide to another. Different routers have to perform different routing tasks and therefore vary in the number and type of μ Rings. However, there are two basic router designs in Amon that apply to most of the possible routers in the network. For illustration, we show switches 33 and 38 in Fig. 2 and 4, which reside in the south-west submesh ($Mesh_{SW}$) of the 64-node Amon. We chose these two switches as representative of where and how μ Ring filters are placed in Amon to implement the routing strategy explained above. In principle, μ Ring filters switch according optical signals to the nodes in the same column or row, based on where the data came from and what the relative position of the switch is.

Fig. 2 and 4 depict the switch architecture (left) and supported turns (right). Red and black links are coloured in the same way as in the layout in Fig. 5 to indicate when logically different links such as Intermesh links and submesh links, are physically the same.

Switch33: Switching has to be provided when the optical signal has to be dropped into another waveguide: (i): Data is coming in on the Intermesh link from $Mesh_{NW}$. Possible destinations of the data could either be nodes in the column below N_{33} or in any of the columns to the right. In the latter case, no switching is needed and the optical data stream just passes through and will be switched in one of the subsequent nodes in this row. In the former case, the optical stream carrying the data has to be filtered from the current waveguide to the southwards waveguide in this column. This is carried out by the μ Ring filter bank **S1** in Fig. 2, which has to be tuned to all the wavelength sets in the column below. In case of a 64-node Amon, this would be λ_5 , λ_9 , and λ_{13} , for addressing N_{37} , N_{41} , and N_{45} , respectively (see Fig. 3). This scheme applies to each link in Amon: Data coming in from the left or right must be switched

up or down in the mesh to steer the optical streams to the destinations in the same column.

(ii): Data is coming in on the Intermesh link from $Mesh_{NE}$. The same as in the previous case is applied: As the Intermesh link is physically the same as the mesh link from the east to west, switching has to be provided to steer the optical data to the mesh link from north to south in case the data's destination resides in the same column. This is performed by **S2**, which responds to the same wavelength sets as **S1** since they enable switching to the same destinations, just from other directions.

(iii): Data is coming in on the Intermesh link from $Mesh_{SE}$. This link is physically the same as the (logical) mesh link from the south to north, as illustrated in Fig. 5. If the destination is in the same row as N_{33} , switching has to be provided to the waveguide from west to east. This switching is performed by the μ Ring filter bank **S3**. **S3** therefore provides μ Ring filters for each wavelength set of the nodes to the right, i.e., λ_2 , λ_3 , and λ_4 , for addressing N_{34} , N_{35} , and N_{36} , respectively (Fig. 3).

Besides switching rings, ejection filters (**Ej** in Fig. 2) have to be implemented to receive the data addressed to N_{33} (modulated on λ -set λ_1 in 64-node Amon). Data could potentially be sent in on the mesh links to the east and south, as well as on the Intermesh link to the west. In addition, injection rings (**Inj**) are placed on all waveguides to modulate new data in the corresponding local or Intermesh link.

Switch38: The following cases have to be considered: (i): Data is coming in from the north in the submesh. If the destination of the data is in the same row as Sw_{38} , it has to provide μ Ring filters to switch the data to the left or right, based on the wavelength set of the incoming stream. This functionality is implemented by **S4**. **S4** thus corresponds to the wavelength sets of the destination to the left and right.

(ii): Data is coming in from the south and is meant for a node in the same row. This functionality is implemented by **S3**, which responds to the same wavelength sets as **S4**, as they switch data to the same destinations but from a different waveguide. The link coming from the south is physically the same as the Intermesh link from $Mesh_{SE}$.

(iii): Data is coming in from the west and is meant for a destination in the same column. The same switching functionality as in Sw_{33} is applied here, the only difference is that destinations could be both to the south and north, given the central position of Sw_{38} . The **S1** μ Ring filter banks implement this switching to the south and north, respectively.

(iv): Data is coming in from the east and is meant for a destination in the same column. Basically the same applies as in (iii): μ Ring filter bank **S2** tuned to the corresponding wavelength sets are placed accordingly to implement the switching to the north and south.

Similarly to Sw_{33} , ejection filters are placed on all waveguides on which data streams could be sent in. In the case of Sw_{38} , this is from each cardinal direction, leading to four ejection rings. Injection rings are pro-





Fig. 1: 48-node Amon with 4×3 submeshesFig. 2: Switching in Switch33 (in $Mesh_{SW}$ of 64-Node Amon)

Fig. 3: Routing Examples in a 64-Node Amon

Fig. 4: Switching in Switch38 (in $Mesh_{SW}$ of 64-Node Amon)

vided in all directions required by the routing algorithm.

Every other switch is implemented similar to these switch designs, as they all perform the same routing functionality (row/column allocation). If a switch does not need to implement switching to all cardinal directions (e.g. Sw_{33}), corresponding rings are omitted.

3.1.4 VLSI Layout of the Data Network

This section presents our layout that takes laser placement and sharing into consideration, and shows that the logical topology of Amon not only features a low number of μ Ring filters, but also allows an efficient physical implementation. Fig. 5 illustrates how laser sources enable optical communication to the submesh located in the south-west $Mesh_{SW}$. For nodes in every other submesh to send data to any node in $Mesh_{SW}$, they need a laser source to drive the Intermesh links (see Fig. 5). For that purpose, we place the laser sources for the Intermesh links in the corner close to the injection channels, as shown in green in Fig. 5. Intermesh links, colored in red, will physically become a mesh link in the logical topology. This allows a simpler layout with fewer μ Ring filters and laser sources, compared to our initial proposal [4], i.e. having a separate waveguide for every single row/column in the submeshes. A separate waveguide and laser source for the submesh links would decrease the loss and thus power dissipated by the laser source, but also require more laser sources and waveguides. In our layout, $Mesh_{SW}$ only needs one more laser source to power the waveguides in the submesh from north to south. Current Ge-on-Si multimode lasers with a gain spectrum of 1590-1610nm and mode spacing of

0.063nm have been demonstrated [30] and could provide around 300 wavelengths per laser source. Therefore, as a laser source in Amon has to provide $N/4$ wavelengths, one laser source (and waveguide) for all Intermesh links in one direction of a submesh is sufficient for moderate network sizes and bandwidth requirements. Similarly to the layout illustrated in Fig. 5, laser sources and waveguides are placed in every submesh. This allows an efficient routing of waveguides with a low number of waveguide crossings.

Laser source sharing: The layout proposed in Fig. 6 has the advantage of placing the laser sources in close proximity to the according sending nodes, allowing short paths and low path losses. Given that all laser sources output the same wavelengths, there is further potential in our design for laser source sharing. Recent studies showed that sharing a laser source over up to 16 waveguides is beneficial in terms of laser power [31], before splitter loss starts to neutralize the benefits again. Moreover, placing laser sources on the edges of the chip—where temperature fluctuations are less frequent [32]—decreases the impact of temperature on the laser efficiency and, in turn, laser power, and does not constitute a major imposition on path loss at current waveguide loss parameters of 0.1dB/mm [21]. This facilitates chip packaging as fewer laser sources have to be coupled into the chip. We therefore propose another sharing and placement of the laser source as shown in Fig. 7, in which we combine the designated laser sources and place them at the edges of the chip. Laser sources that are within the bounds of the same blue box in the left




Fig. 5: Layout and Laser Source Placement for the Lower-Left Submesh




Fig. 6: Complete layout




Fig. 7: Complete layout: 8 vs. 4 laser sources

figure are combined to one laser source in accordance to the numbering. Next to the benefit of having the laser sources exposed to fewer temperature fluctuations, only half the number of laser sources are required.

Fig. 8 shows the maximum IL path for 8 and 4 laser source design, which is closely related to the required laser output power. The longest path remains the same for both layouts, namely from N_{20} to N_{45} . The maximum insertion loss path increases by only 12% when the illustrated laser source sharing is applied. In return, however, only four laser sources are required as opposed to eight. Our results obtained by DSENT [33] show that this sharing of laser sources allows to save 15% of laser power in the 64-node Amon.

The straight-forward, tile-based layout of Amon decreases waveguide bending, and allows optical signals to be dropped at most once to reach any of its destination (excluding the ejection filter). Waveguide crossings also contribute to the total IL; however, 3D integration can be utilized to minimize its effects [21]. All these attributes of Amon's layout outline its effectiveness regarding IL and laser power.

3.2 Control Network

The CN should introduce as little overhead as possible in terms of latency, energy, and area, while successfully managing destination-checking. Implementing the CN optically rather than electrically was shown to be superior in both latency and energy, particularly when a global CN is required [34] [29]. The destination to be checked might be far away on the chip, which would introduce larger number of hops and in turn large latency and energy overheads, thereby diminishing the benefits of an optical data network. An optical CN, on

the other hand, can leverage the distance-independent energy consumption and signal propagation of light to reach all destination efficiently, making it ideal for a global CN. However, to avoid large static optical power overheads, its bandwidth should be kept low, thus requiring small control packets for low-latency destination-checking, which we target in our proposal.

Nodes exchange the following control packets prior to data transmission: First, a REQ is sent to the node it wants to send data to. Once the REQ has been received by the destination node, it checks whether it is free to receive data. If so, it sends an ACK to the requesting node indicating that it can start sending its data packet. As soon as a REQ is received, the sender starts transmitting its data. If the destination is busy, it will keep the REQ in its buffers and send out the ACK once it is free again.

In Amon's CN (Fig. 9), every node sends its REQs/ACKs on one distinct wavelength and can transmit on every waveguide, while it can receive only on one waveguide. Groups of 8 nodes share one waveguide for receiving, where the optical stream is split to all receivers. Assuming $N = 64$, this makes 8 distinct waveguides for the CN in our design. If N_0 wants to send a REQ to N_{63} , it modulates its REQ on waveguide 7. This REQ will be received by every destination connected to this waveguide, therefore the requester has to encode the destination ID into the packets which will be checked by the receivers. Similarly, destinations modulate their ACKs on the waveguide where the sender node is attached to, and encode their source ID for the sender to identify from which destination the packet came from. In addition to the node IDs, the packet type must be encoded to distinguish REQs from ACKs. We improve upon the CN proposed in [29] in the following ways: [29] uses the same principle as shown in Fig. 9, but with 16 nodes attached to four waveguides. 8 nodes per waveguide requires 3 bits for ID encoding. We support two packet types (REQ/ACK): In case a destination is busy, it will store the REQs and send the according ACKs out once it is free again. This leads to a packet size of 4 bits. Keeping all alive REQs at a destination requires buffering. However, REQs are only 4 bit long and there will be a maximum of one per issuing nodes, thus requiring only very little buffer space ($N \times 4$).

The CN does not require switching, making its layout




Fig. 8: Maximum Insertion Loss (IL) Path for different Laser Source Sharing




Fig. 9: Control Network Design




Fig. 10: CN Layout




Fig. 11: Amon: Die Layers

straight-forward: Since every node in the NoC has to be passed to interface the CN, we implement it as shown in Fig. 10. A light is coupled into the waveguides, which form a loop that allows every node to address every other node. To avoid waveguide crossings and facilitate the physical implementation, Amon assumes a 3D integrated chip and places the CN on a separate optical layer, as depicted in Fig. 11. For more information on interconnecting multiple die layers with Through-silicon Vias (TSVs), we refer the reader to [21] [19].

4 EVALUATION

4.1 Experimental Set-Up

In this section, we compare Amon to some state-of-the-art NoCs: QuT [29](all-optical), Atac [23], Firefly [12], Meteor [22](hybrid), and a baseline electrical 2D mesh. We use DSENT [33] for area, energy, and power estimations with the IL parameters in Table 1, and Sniper [35] for performance modeling of SPLASH-2 [36] and PARSEC [37] applications with the *sim-large* input set. Dynamic power during these applications is obtained by inserting energy values obtained from DSENT into the Sniper simulation infrastructure. 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 [38]. We use a 22nm technology, 5Ghz core/router 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 [39], and assume a data packet size of 256 bits and flit size of 64 bit. Electrical

links in all NoCs are 64-bit wide. We study the NoC designs for 64 nodes and assume an 8×8 core layout.

QuT [29] is a low power WRONoC that uses passive μ Rings 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. As every destination has one ejection channel, a separate CN, like in Amon, is required for contention-resolution. QuT's CN is similar to Amon's, but has only 4 waveguides with 16 destinations each. Busy destinations send out NACK packets, upon which senders will wait for a back-off time and retransmit their REQ. Control packets are modulated on one wavelength, and data packets on 8 wavelengths [29].

Atac [23] consists of a 2D electrical mesh, overlaid by an optical network (ONET). In the 64-node version, each node is connected to the ONET—a bundle of 32 Single-Writer-Multiple-Reader links [34] with 64 wavelengths each. Packets for destinations closer than 4 hops away are sent on the electrical mesh; others on the ONET.

Firefly [12] 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 within the same cluster, and optical links to their duals for inter-cluster communication. Prior to data transmission on optical links, control packets are exchanged and rings tuned/detuned to reduce laser power. We assume enough 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.

Meteor [22], similar to Atac, implements a 2D electrical mesh overlaid with an ONET. However, 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 MWMR buses. Each hub can send on a 128-bit wide link.

2D Mesh is a conventional 2D electrical mesh network with XY-routing and wormhole switching. Routers can take up to 5 pipeline stages when clocked at 5Ghz, such as in Intel's TeraFLOPS design [2]. 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




Fig. 12: Average packet latency on synthetic workloads




Fig. 13: Average Packet latency on PARSEC/SPLASH-2 workloads

is the *de facto* standard in industry [40] and to outline the benefits of implementing optical links.

Amon is implemented as described in the previous section. Packets are modulated on 8 wavelengths on the data network, and on 1 wavelength on the CN, like QuT for a meaningful comparison.

4.2 Packet Latency

Fig. 12 shows the latency for traditional synthetic traffic patterns. In Hotspot traffic, 20% of the nodes, placed in the top rows of the tiled chip layout, receive 80% of the total traffic, while the rest is distributed randomly across all nodes. Latency on the optical path includes delay in the E/O and O/E backends (23.8ps and 4.2ps, respectively [29]), signal propagation on the waveguide (11ps/mm [29]), as well as modulation delay required to serialize packets exceeding the link bandwidth. In Amon, the CN and data modulation are the major contributors to packet latency, as light propagates fast and optical bandwidth is fairly low. Compared to QuT, Amon slightly improves latency and throughput thanks to a more efficient CN: In QuT, grouping 16 nodes per waveguide requires 4 bits for ID encoding, and 3 packet types (REQ/ACK/NACK) another 2 bits - 6 bits in total. With 1λ , this requires one more cycle compared to Amon's control packets, which are of 4 bit size. Overall, both of these networks show a constant performance for all traffic patterns apart from Hotspot, which causes high contention at the destination nodes. Compared to other NoCs, both Amon and QuT offer significantly better throughput than the alternative NoCs, apart from Neighbor and Hotspot traffic. In the former, Atac, Meteor

and 2D Mesh are superior as they have electrical mesh links which are more efficient for local traffic.

Fig. 13 shows the average packet latency normalized to Amon over a range of PARSEC/SPLASH-2 applications. Similar to Hotspot traffic, Amon and QuT perform weak in applications where there is high contention, i.e. where a high number of nodes try to access the same destination. In most PARSEC and SPLASH-2 applications, however, the average injection rate over the entire execution time is fairly low [43], alleviating the drawbacks of a CN. As with synthetic traffic, Amon decreases the packet latency slightly compared to QuT. Apart from 2D Mesh, the remaining alternative NoCs provide slightly less packet latency than Amon. This is mainly due to the large serialization delay required for cache line transfers, which are 576-bit long in the used CPU architecture: With 8λ , each sender needs 36 cycles to modulate these large data packets. This is only beneficial in Amon when the distance between source and destination is large.

4.3 Power Consumption

Advances in both architectures and technology are imperative for the adoption of ONoCs. To assess the potential of improvements in device technologies, we evaluate power consumption under both conservative and aggressive technology parameters. Power results for *conservative* and *aggressive* IL and heating parameters (see Table 1) are depicted in Fig. 14 and 15, respectively.

Amon needs fewer μ Rings for switching and provides a more efficient layout than QuT, leading to lower μ Ring heater and laser power. In the conservative case, Amon reduces power by 30% compared to QuT. The global optical crossbar in Atac entails large power overheads, leading to more than $3\times$ the power of Amon. Meteor and Firefly are hybrid NoCs relying on both electrical and optical links, thereby reducing the amount of optical resources in the NoC. While they decrease laser and μ Ring heater power, they consume more power in the electrical routers and links. While the power savings in the conservative case compared to hybrid and electrical NoCs is slim, aggressive technology parameters widen up this gap significantly, with Amon saving at least 50% compared to hybrid and electrical NoCs. Compared to QuT, the benefits are reduced, as the excess of optical resources needed by QuT comes at a lower cost in terms of power. However, Amon still saves 24%. Amon implemented with conservative optical parameters almost doubles the power consumption compared to the aggressive case, outlining the benefits of technology advances in all-optical NoCs. The low bandwidth required on the control network leads to little laser and μ Ring heater power, making the CN a minor factor with regards to the total power consumption contribution: $\sim 18\%$. Dynamic power is generally very low in ONoCs due to very low energy for transmitting a bit (E/O and O/E consume just 100fJ/bit [29]).

We study the relation of the maximum achieved throughput under uniform traffic and its required

TABLE 1: Insertion Loss Parameters

Parameter	Conservative	Aggressive
Laser efficiency	0.25 [19]	0.3 [20]
Coupler	2 dB [20]	1 dB [19]
Ring; Through	0.001 dB [21]	0.0001 dB [20]
Ring; Drop	1.5 dB [21]	1 dB [33]
Waveguide Bending	0.005 dB [29]	0.005 dB [29]
Modulator Insertion	0.001 dB [21]	0.001 dB [21]
Waveguide prop.	0.2 dB/mm [20]	0.1 dB/mm [41]
Waveguide crossing	0.12 dB [29]	0.05 dB [21]
Splitter	0.2 dB [21]	0.1 dB [29]
Photodetector loss	1 dB [33]	0.1 dB [22]
μ Ring heater power	$20\mu\text{W}/\mu\text{Ring}$ [28]	$5\mu\text{W}/\mu\text{Rings}$ [42]

power before network saturation by computing the Throughput-per-Watt (TPW), depicted in Fig. 16 and Fig. 17 for conservative and aggressive optical technology parameters, respectively.

Amon improves the TPW of QuT by 23% due to lower power consumption. Meteor is the most competitive alternative NoC with 78% of Amon's TPW. All other NoCs provide less than half of that. These benefits become more eminent with aggressive technology parameters as they allow to almost double the TPW compared to conservative ones. Although the TPW benefits of Amon compared to QuT are slightly less than in the conservative case (due to less power benefits, see previous paragraph), TPW is improved by between 31% and 75%.

To assess the power efficiency for realistic traffic we compute the Power-Delay-Product (PDP), shown in Fig. 18 and Fig. 19. The power efficiency of Amon with conservative parameters improves, on average, by 24%, 210%, and 50% compared to QuT, Atac, and 2D Mesh, respectively. Meteor and Firefly have a similar PDP as Amon, mainly due to lower dynamic power in their electrical network caused by low injection rates of PARSEC/SPLASH-2 applications. For aggressive optical parameters, however, Amon's PDP improves significantly (90%), leading to higher power-efficiency compared to all other NoCs.

4.4 Area

Tab. 2 lists the area normalized to Amon. The results include area for routers and electrical links, as well as waveguides ($4\mu\text{m}$ pitch [33]), μ Rings ($5\mu\text{m}$ diameter [21]) and E/O and O/E backend circuitry. Optical components require more area than electrical. This is reflected in the area results when one compares all-optical to hybrid or electrical NoCs. The hybrid NoCs Firefly and Meteor needs 34% and 24% less area than Amon. Atac has a large area overhead due to its global crossbar. The 2D Mesh reduces area by 39%, but pays large power and latency overheads. Although Amon needs 8 waveguides on the CN compared to QuT which needs only 4, Amon has fewer μ Rings and provides a more straight-forward, efficient layout, leading to 21% area savings.

4.5 Scalability

In all-optical NoCs, the number of μ Rings is an important metric to assess scalability since it not only directly relates to μ Ring heating power, but also adds to the path losses. The power breakdowns in Fig. 14/ 15 depict this, with more than half of the power budget




Fig. 14: Power Breakdown (Conservative)




Fig. 15: Power Breakdown (Aggressive)

TABLE 2: Area normalized to Amon

	Amon	QuT	Atac	Firefly	Meteor	2D Mesh
Normalized Area	1.0	1.21	1.73	0.66	0.76	0.61

TABLE 3: Number of μ Rings for varying NoC sizes

Nodes	Amon	QuT	Control Network
64	36235	45056	1024
144	170448	188654	5184
256	544302	671744	16384

being utilized for μ Ring heating. We list the number of μ Rings for the data and control network (of Amon) for different network sizes for Amon and QuT Table 3. We chose this comparison since QuT is all-optical and scales without clustering, while the other considered NoCs all propose network scaling by utilizing some sort of clustering [22] [23] [12].

Amon requires fewer μ Rings at all network sizes. Apart from μ Ring heater power, Amon, given its mesh-based layout, also provides a simpler topology to scale than QuT, which is a ring topology with cross links, which might introduce a challenging network layout, potentially leading to additional path losses. The number of wavelengths on the data network in both NoCs scales equally ($(N/4) \times 8$ for N nodes). Although providing better scalability than QuT, the number of μ Rings in Amon scales as $O(N^2)$, which leads to considerable power requirements which would limit its scalability when compared to hybrid or electrical NoCs. Because of this quadratic nature it is widely acknowledged that WRONoCs may not be able to scale over hundreds of endpoints. Core clustering, in which ONoC routers are shared among nodes, would therefore be a preferable approach for scaling up Amon. Indeed, it would be even possible to use Amon as the ONET of a hybrid NoC similar to ATAC or Meteor.

The CN becomes less significant as network size increases since only one wavelength is assigned to each node for modulation, as opposed to eight on the data network. However, the numbers of waveguides required on the CN might become difficult to implement, since it scales with $N/8$ (see Section 3.2). Higher splitting degrees to decrease the number of waveguides alleviates this issue, however, also leads to more μ Rings, longer control packets for ID encoding and, in turn, to higher latency and energy on the CN. Scalability of the CN is thus an interesting study for future work.




Fig. 16: Throughput-per-Watt (Conservative)




Fig. 18: Power-Delay-Product (Conservative)

4.6 Discussion

Our experiments have shown that Amon improves upon the state-of-the-art all-optical NoC QuT by providing a more efficient topology, layout, and arbitration mechanism. Furthermore we have found that all-optical NoCs can be designed to be competitive alternatives to electrical or hybrid NoCs, especially at moderate to high injection rates, as dynamic power is much lower than in its counterparts. Contention on Amon's CN may lead to latency overheads for pathological traffic patterns at high injection rates. However, Amon uses optical data transmission only, which provides low-latency data transmission. This is particularly beneficial for traffic patterns that contain data communication between distant nodes on the chip. For local traffic, electrical links in hybrid NoCs are more effective as they avoid E/O and O/E overheads.

Being all-optical, Amon has large static power requirements. Even with our efficient layout and decreased number of μ Rings, the overheads are considerable and makes it hard to provide latency benefits large enough to outweigh these power overheads, compared to hybrid or electrical NoCs for low injection rates. However, silicon photonics are constantly improving, and advanced devices allow for less insertion loss, higher laser efficiencies, and lower μ Ring heating. Once the latest device technologies are in place for widespread manufacturing, our experiments suggest that all-optical NoCs will soon dominate their electrical counterparts. Designs like Amon that make efficient use of these components will then guide designers to make efficient use of optical on-chip communication.

5 CONCLUSION

This paper outlines that advances in nanophotonic device technology and optical Network-on-Chip architectures make optical on-chip interconnects increasingly superior compared to their electrical counterparts. Our novel ONoC design Amon improves upon the state-of-the-art by decreasing μ Ring heater and laser power by




Fig. 17: Throughput-per-Watt (Aggressive)




Fig. 19: Power-Delay-Product (Aggressive)

providing an efficient topology for wavelength routing. We outlined the benefits and trade-offs of Amon compared to a wide range of optical, hybrid, and electrical NoCs, and exposed the potential of aggressive device parameters. Amon also provides interesting future studies dealing with adaptive laser sources [44]: Destination-checking in Amon requires the exchange of control packets. In addition to access control, these packets could be used to control an adaptive laser source, which would decrease laser power in a latency-efficient manner.

ACKNOWLEDGEMENTS

This work has been funded by the Engineering and Physical Sciences Research Council (EPSRC) through grant EP/K015699/1 and by the European Unions Horizon 2020 research and innovation programme under grant agreement No 671553.

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 *et al.*, "An 80-tile 1.28 tflops network-on-chip in 65nm cmos," in *ISSCC*. IEEE, 2007, pp. 98–589.
- [3] T. Corporation, "Tilera multicore processors," <http://www.tilera.com/products/processors>, 2007.
- [4] S. Werner, J. Navaridas, and M. Luján, "Amon: An advanced mesh-like optical noc," in *HOTI*. IEEE, 2015, pp. 52–59.
- [5] J. Chan and K. Bergman, "Photonic interconnection network architectures using wavelength-selective spatial routing for chip-scale communications," *JOCN*, vol. 4, no. 3, pp. 189–201, 2012.
- [6] I. Artundo, W. Heirman, M. Loperena, C. Debaes, J. Van Campenhout, and H. Thienpont, "Low-power reconfigurable network architecture for on-chip photonic interconnects," in *HOTI*. IEEE, 2009, pp. 163–169.
- [7] J. Chan, G. Hendry, A. Biberman, and K. Bergman, "Architectural exploration of chip-scale photonic interconnection network designs using physical-layer analysis," *Journal of Lightwave Technology*, vol. 28, no. 9, pp. 1305–1315, 2010.
- [8] G. Hendry, S. Kamil, A. Biberman, J. Chan, B. G. Lee, M. Mohiyuddin, A. Jain, K. Bergman, L. P. Carloni, J. Kubiatowicz *et al.*, "Analysis of photonic networks for a chip multiprocessor using scientific applications," in *NOCS*. IEEE, 2009, pp. 104–113.
- [9] N. Kirman, M. Kirman, R. K. Dokania, J. F. Martinez, A. B. Apsel, M. A. Watkins, and D. H. Albonesi, "Leveraging optical technology in future bus-based chip multiprocessors," in *MICRO*. IEEE Computer Society, 2006, pp. 492–503.

- [10] S. Pasricha and N. Dutt, "Orb: An on-chip optical ring bus communication architecture for multi-processor systems-on-chip," in *ASP-DAC*. IEEE Computer Society Press, 2008, pp. 789–794.
- [11] M. J. Cianchetti, J. C. Kerekes, and D. H. Albonesi, "Phastlane: a rapid transit optical routing network," *ACM SIGARCH Computer Architecture News*, vol. 37, no. 3, pp. 441–450, 2009.
- [12] 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, no. 3. ACM, 2009, pp. 429–440.
- [13] A. Joshi, C. Batten, Y.-J. Kwon, S. Beamer, I. Shamim, K. Asanovic, and V. Stojanovic, "Silicon-photonic cros networks for global on-chip communication," in *NOCS*. IEEE Computer Society, 2009, pp. 124–133.
- [14] Y. Pan, J. Kim, and G. Memik, "Flexishare: Channel sharing for an energy-efficient nanophotonic crossbar," in *HPCA*. IEEE, 2010, pp. 1–12.
- [15] 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.
- [16] X. Zhang and A. Louri, "A multilayer nanophotonic interconnection network for on-chip many-core communications," in *DAC*. ACM, 2010, pp. 156–161.
- [17] Y.-H. Kao and H. J. Chao, "Blocon: A bufferless photonic cros network-on-chip architecture," in *NOCS*. IEEE, 2011, pp. 81–88.
- [18] 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," *ACM Journal on Emerging Technologies in Computing Systems*, vol. 7, no. 2, p. 8, 2011.
- [19] 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*. IEEE Computer Society, 2012, pp. 282–293.
- [20] 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 *et al.*, "Cross-layer energy and performance evaluation of a nanophotonic manycore processor system using real application workloads," in *IPDPS*. IEEE, 2012, pp. 1117–1130.
- [21] C. Li, M. Browning, P. V. Gratz, and S. Palermo, "Luminoc: A power-efficient, high-performance, photonic network-on-chip," *TCAD*, vol. 33, no. 6, pp. 826–838, 2014.
- [22] S. Bahirat and S. Pasricha, "Meteor: Hybrid photonic ring-mesh network-on-chip for multicore architectures," *TECS*, vol. 13, no. 3s, p. 116, 2014.
- [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*. ACM, 2010, pp. 477–488.
- [24] M. Briere, B. Girodias, Y. Bouchebaba, G. Nicolescu, F. Mieyeville, F. Gaffiot, and I. O'Connor, "System level assessment of an optical noc in an mpoc platform," in *DATE*. EDA Consortium, 2007, pp. 1084–1089.
- [25] 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, no. 3. IEEE Computer Society, 2008, pp. 153–164.
- [26] 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*. IEEE, 2011, pp. 1–6.
- [27] 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*. EDA Consortium, 2013, pp. 1589–1594.
- [28] S. Koohi and S. Hessabi, "Scalable architecture for a contention-free optical network on-chip," *Journal of Parallel and Distributed Computing*, vol. 72, no. 11, pp. 1493–1506, 2012.
- [29] P. K. Hamedani, N. E. Jerger, and S. Hessabi, "Qut: A low-power optical network-on-chip," in *NOCS*. IEEE, 2014, pp. 80–87.
- [30] J. Liu, X. Sun, R. Camacho-Aguilera, L. C. Kimerling, and J. Michel, "Ge-on-si laser operating at room temperature," *Optics letters*, vol. 35, no. 5, pp. 679–681, 2010.
- [31] C. Chen, T. Zhang, P. Contu, J. Klamkin, A. K. Coskun, and A. Joshi, "Sharing and placement of on-chip laser sources in silicon-photonic nocs," in *NOCS*. IEEE, 2014, pp. 88–95.
- [32] M. P. Gupta, M. Cho, S. Mukhopadhyay, and S. Kumar, "Thermal investigation into power multiplexing for homogeneous many-core processors," *Journal of Heat Transfer*, vol. 134, no. 6, p. 061401, 2012.
- [33] C. Sun, C.-H. O. Chen, G. Kurian, L. Wei, J. Miller, A. Agarwal, L.-S. Peh, and V. Stojanovic, "Dscent-a tool connecting emerging photonics with electronics for opto-electronic networks-on-chip modeling," in *NOCS*. IEEE, 2012, pp. 201–210.
- [34] K. Bergman, L. P. Carloni, A. Biberman, J. Chan, and G. Hendry, *Photonic network-on-chip design*. Springer, 2014.
- [35] W. Heirman, T. Carlson, and L. Eeckhout, "Sniper: scalable and accurate parallel multi-core simulation," in *ACACES*. HiPEAC, 2012, pp. 91–94.
- [36] S. C. Woo, M. Ohara, E. Torrie, J. P. Singh, and A. Gupta, "The splash-2 programs: Characterization and methodological considerations," in *ISCA*. IEEE, 1995, pp. 24–36.
- [37] C. Bienia, S. Kumar, J. P. Singh, and K. Li, "The parsec benchmark suite: Characterization and architectural implications," in *PACT*. ACM, 2008, pp. 72–81.
- [38] Intel, "Intel Xeon Processor 5500 Series," Tech. Rep., 2009. [Online]. Available: <http://ark.intel.com/products/37106>
- [39] Y. Ben-Itzhak, E. Zahavi, I. Cidon, and A. Kolodny, "Hnocs: Modular open-source simulator for heterogeneous nocs," in *SAMOS*. IEEE, 2012, pp. 51–57.
- [40] C. Ramey, "Tile-gx100 manycore processor: Acceleration interfaces and architecture," in *HC*, 2011.
- [41] M. Lipson, "Guiding, modulating, and emitting light on silicon-challenges and opportunities," *Journal of Lightwave Technology*, vol. 23, no. 12, p. 4222, 2005.
- [42] Y. Demir and N. Hardavellas, "Parka: Thermally insulated nanophotonic interconnects," in *NOCS*. ACM, 2015, p. 1.
- [43] J. Lee, C. Nicopoulos, S. J. Park, M. Swaminathan, and J. Kim, "Do we need wide flits in networks-on-chip?" in *ISVLSI*. IEEE, 2013, pp. 2–7.
- [44] Y. Demir and N. Hardavellas, "Ecolaser: an adaptive laser control for energy-efficient on-chip photonic interconnects," in *ISLPED*. ACM, 2014, pp. 3–8.



Sebastian Werner is a PhD student in computer architecture in the Advanced Processors Technologies group of the University of Manchester. Sebastian received the B.S. and M.S. degree in Computer Engineering from the Darmstadt University of Technology in 2013 and 2014, respectively. His current research interests include scalable and energy-efficient networks-on-chip, silicon photonics, and many-core architectures.



Dr. Javier Navaridas is a Lecturer in computer architecture in the University of Manchester. Javier obtained his MEng in Computer Engineering in 2005 and his PhD in Computer Engineering (awarded with an Extraordinary Doctorate Award - top 5% theses) in 2009, both from the University of the Basque Country, Spain. His research interests encompasses interconnects, parallel and distributed systems, computer architecture, performance evaluation and characterization of application's behaviour.



Dr. Mikel Luján is a Royal Society University Research Fellow in the School of Computer Science at the University of Manchester. His research interests include managed runtime environments and virtualization, many-core architectures, and application-specific systems and optimizations. Luján has a Ph.D. in computer science from the University of Manchester.