

Reconfigurable Optical and Wireless (R-OWN) Network-on-Chip for High Performance Computing

Md Ashif I Sikder
School of Electrical
Engineering and Computer
Science, Ohio University
Athens, OH-45701
ms047914@ohio.edu

Avinash K Kodi
School of Electrical
Engineering and Computer
Science, Ohio University
Athens, OH-45701
kodi@ohio.edu

Ahmed Louri
Department of Electrical and
Computer Engineering,
George Washington University
Washington, DC-20052
louri@gwu.edu

ABSTRACT

With the scaling of technology, the industry is experiencing a shift from multi-core to many-core architectures. However, traditional on-chip metallic interconnects may not scale to support these many-core architectures due to the increased hop count, high power dissipation, and increased latency. Recently, attention has recently been shifted to emerging technologies such as optical and wireless interconnects for future on-chip communications. Although emerging technologies show promising results for power-efficient, low-latency, and scalable on-chip interconnects, the use of single technology may not be sufficient to support future many-core architectures. In this paper, we propose a Reconfigurable Optical-Wireless Network-on-Chip (R-OWN) that facilitates communication through static optical links and reconfigurable wireless links. The network diameter of R-OWN is restricted to three hops by dividing the network into several optical domains of 64-cores (called a cluster) and by connecting the clusters using one-hop wireless network. The optical bandwidth is efficiently shared using time division multiplexing (TDM), and the wireless bandwidth is shared using frequency division multiplexing (FDM). Packets routed across optical and wireless networks are proved to be deadlock-free. Our results indicate that R-OWN improves energy-efficiency by 44-51%, performance (throughput and latency) by 13-31%, and area by 4-13% when compared to state-of-the-art wired, wireless, optical, and hybrid on-chip networks.

CCS Concepts

•Computer systems organization → Embedded systems; *Redundancy*; •Networks → Network reliability;

Keywords

Network-on-Chip, Wireless, Optical, Reconfiguration, Antennas, Transceivers.

1. INTRODUCTION

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

NANOCOM'16, September 28-30, 2016, New York, NY, USA

© 2016 ACM. ISBN 978-1-4503-4061-8/16/09...\$15.00

DOI: <http://dx.doi.org/10.1145/2967446.2967457>

The number of processing cores on a chip is increasing with the scaling down of transistor size. Hence, we see a shift from multicore to manycore architecture in computer industry. For example, KALRAY has published a MPPA (Multi-Purpose Processor Array) manycore roadmap that features 256 to 1024 processing cores on a single die [1]. Adapteva, a company that builds energy efficient microprocessors, has claimed to fabricate 1024-core processors [13]. Therefore, the era of kilo-core processors is not far away. However, according to the International Technology Roadmap for Semiconductors (ITRS), the development in metallic interconnection technology would not be sufficient to support these future manycore processors. As a result, researchers are focusing on emerging technologies such as optics, radio-frequency (RF)/wireless, and 3D-NoC for future manycore architectures.

Optical interconnects offer several advantages such as low energy consumption (~ 0.25 pJ/bit), reduced link latency (\sim ps), and increased bandwidth (~ 40 Gbps) via wavelength division multiplexing (WDM)—all of which makes optics a suitable technology for on-chip communications [17], [14], [7], [3]. However, for large core counts, crossbar-based architectures such as Corona [17], Firefly [14], and 3D-NoC [12] suffer from scalability issues because the number of ring modulators and photodetectors increase quadratically with the number of cores, thus resulting in high area and power overhead. Nevertheless, an alternate choice is to use hierarchical optical interconnects with a lower number of cores in each hierarchy and a complex optical wavelength(s) sharing/allocation; but this may increase fabrication complexity and both energy-consumption and network latency due to multiple hops requirement.

Wireless NoCs has the advantages of distance-independent energy requirement, CMOS compatibility, and omnidirectional communication—all of which can decrease the energy consumption with reduced hop count [10], [6]. However, on-chip wireless technology requires a higher energy per bit transmission, and mmwave transceivers at 60 GHz center frequency offer fewer channels that can be implemented in current technology while larger bandwidth transceivers operating at 100-300 GHz are more challenging to implement [11]. Therefore, wireless architectures are designed as a two-tier hybrid architecture and combine time division multiplexing (TDM), frequency division multiplexing (FDM), and space division multiplexing (SDM) to overcome the bandwidth limitation [6], [10], [9].

Optics offers several advantages such as energy-efficiency

and high bandwidth but for flat systems (not hierarchical designs), scaling to large core counts becomes prohibitively expensive due to hardware resources and insertion losses. Similarly, omnidirectional wireless enables ease of communication but limited frequency spectrum and higher energy/bit limits the use of wireless technology. Therefore, we propose an architecture called Reconfigurable Optical-Wireless Network-on-Chip (R-OWN)—to exploit the advantages of the emerging technologies while circumventing the disadvantages. We propose to utilize optics for a smaller number of cores (called a cluster) and utilize wireless technology to connect the clusters to facilitate irregular communication. In this architecture, reconfigurable wireless links are designed to adapt to diverse communication patterns. The advantage is that a fraction of the bandwidth from an under-utilized channel can be assigned to an over-utilized channel which will enhance the network performance. Instead of wireless, a second level optical or metallic interconnects could be used to connect the clusters, but (1) metallic interconnects would not scale for a higher number of cores (say 1024-cores), and reconfiguring a metallic interconnect (for example using power gating) would increase network complexity; (2) multiple optical layers would cause heat dissipation problem that could deteriorate the network performance because optics is sensitive to heat; optical networks would require constant laser power, and turning off certain wavelengths would require off-chip transmission which would incur additional delay; and also, a higher number of modulators and demodulators would be required for reconfiguration which would increase power loss. In contrast, wireless interconnects are ideal for reallocating bandwidth due to the lack of wires and wide frequency spectrum. Since the antennas are on the chip, they can be turned off if necessary. As a result, we use optics to enhance locality by sharing the high-bandwidth one-hop communication and wireless to provide scalable and reconfigurable global bandwidth that can be effectively shared by several routers.

The paper is organized as follows: First, we discuss the work related to our architecture. Next we describe the proposed R-OWN architecture with communication mechanism and deadlock avoidance technique. Finally we evaluate the performance of the proposed architecture by comparing against the state-of-the-art architectures.

2. RELATED WORK

In this section, we discuss the prior research work relevant to our proposed architecture, R-OWN. Traditional metallic interconnects generally followed 2-D Mesh, Concentrated Mesh (CMesh), and Torus topologies [2]. Since metallic interconnects may not scale for large core counts, architectures employing emerging technologies—such as wireless and optics are proposed. One such architecture that employs wireless technology is WCube [10]. WCube extends the CMesh architecture by inserting micro-wireless router for a subnet or group of routers. The inter-subnet communication uses wireless technology whereas the intra-subnet communication uses wired technology. Similarly, iWISE [4] uses wireless technology, and WiNoC [6] uses both wired and wireless technology for inter-subnet communication. iWISE is extended to A-WiNoC [5] by providing reconfigurability in wireless channels to reduce network contention. After a fixed time frame, A-WiNoC assigns the wireless links to subnets which is chosen by the global arbiter based on the

outgoing traffic of all the subnets.

Optical NoCs are drawing considerable interest due to their inherent energy and bandwidth advantages. Corona [17] proposes an optical ring-crossbar network using the broadcasting capability of the optical links. Off-chip laser source and dense wavelength division multiplexing (DWDM) is used for data communication. However, Corona requires a very high number of ring resonators and consumes high power as a portion of the wavelength is peeled off by every router on the path. Firefly [14] reduces the optical crossbar costs by utilizing electrical mesh while 3D-NoC [12] reduces the cost utilizing decomposed crossbars. Similar to 3D-NoC, OWN [15] proposes to use smaller crossbars to reduce the cost but uses wireless technology to connect the crossbars. Compared to several proposed architectures, except OWN that use emerging technologies, R-OWN combines both optics and wireless technologies with reconfigurability within the same architecture. In addition, the primary differences between R-OWN and OWN are: (1) the use of reconfigurable wireless links, (2) inclusion of local arbiter, and wireless link bandwidth optimization, and (3) deadlock avoidance technique with higher input buffer utilization.

3. PROPOSED R-OWN ARCHITECTURE

In this section, first, we briefly explain the 64-core cluster architecture. Second, we describe the design of R-OWN for 256 cores and describe routing mechanisms. Third, we analyze deadlock situations especially when packets flow from multiple domains (optics to wireless) and describe a deadlock-free routing methodology.

3.1 64-Core R-OWN Architecture: Cluster

The structure of a 64-core cluster is shown in Figure 1. Each cluster consists of 16 tiles where each tile consists of four cores, four private L1 caches, a shared L2 cache, and a router. For data transmission, a snakelike optical waveguide(s) is shared between the 16 tiles. Hence, to maintain signal integrity, we choose multiple-write-single-read (MWSR) scheme where each tile is assigned a specific wavelength(s), and the other tiles can write on this wavelength(s). Since a tile waits for the optical token to write on a wavelength, latency is increased. However, because the number of tiles sharing an optical token is less than 16, we believe that this latency would not increase abruptly. The structure of a wireless router and optical communication mechanism are shown in Figure 1. For example, receiving the optical token, router 2 modulates the yellow wavelength to send a message to router 11. In addition to data waveguides, an arbitration waveguide is used to circulate token signals.

3.2 Proposed 256-Core R-OWN Architecture

The 256-core R-OWN architecture consists of four 64-core clusters which is shown in Figure 1. Intra-cluster communication is implemented using optical interconnects as described in the previous section, and inter-cluster communication is facilitated using wireless interconnects. Since we have four clusters, twelve (4P_2) unidirectional channels are required to provide cluster-to-cluster wireless communication. Hence, each cluster requires three frequency channels to talk to the rest of the clusters (horizontal, vertical, and diagonal cluster). As a result, each cluster contains three wireless transceiver antennas to send and receive packets to and from the horizontal, vertical and diagonal cluster. In

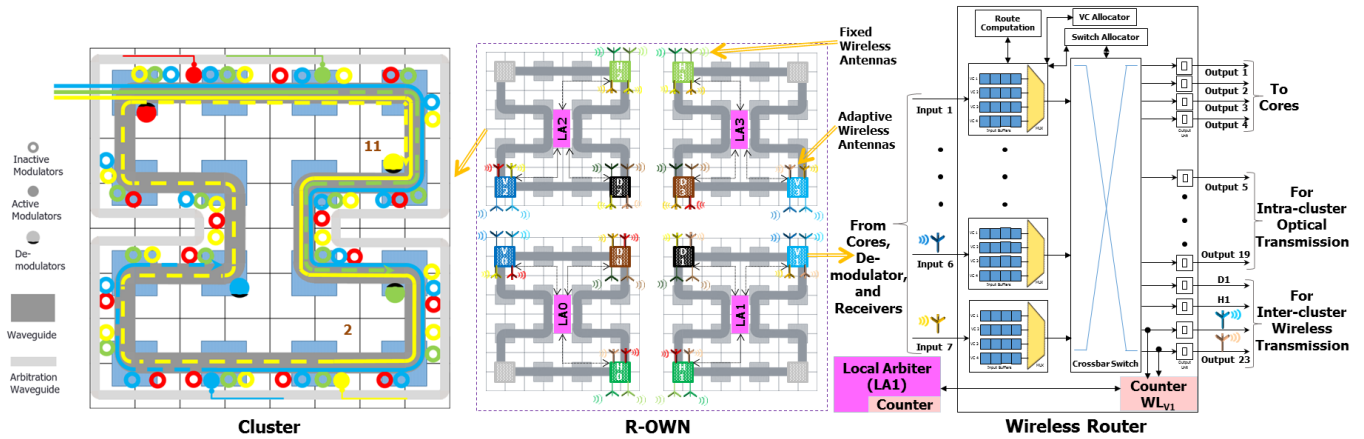


Figure 1: (Left) The 64-Core cluster architecture, connected by a 16×16 optical crossbar. (Middle) 256-Core R-OWN architecture. Routers with the same color communicate with each other, and F_{xy} represent a wireless channel to send packets from cluster x to cluster y . For example, Routers H0 and H1 communicate with each other over frequency channel F01 and F10 respectively while routers V1 and V3 communicate with each other over frequency channel F13 and F31 respectively. (Right) Structure of the proposed wireless router in R-OWN with transmitters, receivers, counters, and local arbiter.

addition to these three fixed wireless channels, each cluster is assigned an adaptive wireless channel. Thus, each wireless router of a cluster contains a transmitter tuned to the adaptive wireless channel frequency assigned to that cluster and a receiver tuned to the adaptive wireless channel frequency assigned to other clusters. However, only one of the three adaptive transceiver antennas can operate for a period of time to maintain signal integrity, and this is ensured by an arbiter (called local arbiter) located inside the cluster. Therefore, a cluster contains three wireless routers, three fixed transceiver antennas to communicate with the horizontal, vertical, and diagonal cluster, three adaptive transceiver antennas, and an arbiter to control the adaptive transceiver antennas. Since we require 16 channels with a total wireless bandwidth of 512 Gbps, the bandwidth of a wireless channel is 32 Gbps.

As shown in Figure 1, three of the four corner routers of each cluster is chosen for the on-chip wireless communication. The corner routers are chosen to provide maximum separation between antennas operating at different frequencies to minimize inter-channel interference. The corner routers (shown in silver color in Figure 1) of 256-core R-OWN are reserved for scaling to 1024-core architecture.

3.3 Wireless Channel Reconfiguration

The adaptive wireless channel of each cluster is reconfigured after a reconfiguration window (set to 100 cycles in our simulation) depending on the number of packets sent to the other clusters. After every 100 cycles, the local arbiter requests for the wireless link usages from the wireless routers of this cluster. Upon receiving the request signal, each wireless router sends its wireless link utilization to the requesting local arbiter. The local arbiter determines the destination cluster of the adaptive wireless link for the next 100 cycles based on the maximum link utilization, resets its counter to zero, sends a decision signal to each of the wireless router of this cluster, and waits for 100 cycles to send again a request signal. Upon receiving the decision signal, a wireless router resets its counter and turns on/off its adaptive antennas. Hence, each wireless router and each arbiter

requires a counter respectively to keep track of the wireless link traversals and to count the number of cycles.

Reconfigurable-Wireless Algorithm	
Step1	Wait for the reconfiguration window, R_W
Step2	Local arbiter, LA_i requests the wireless routers (H_i, V_i, D_i) for wireless link usage ($WL_{H_i}, WL_{V_i}, WL_{D_i}$) where i is the cluster number
Step3	$H_i, V_i,$ and D_i sends $WL_{H_i}, WL_{V_i}, WL_{D_i}$ respectively to LA_i
Step4	LA_i finds the maximum of $[WL_{H_i}, WL_{V_i}, WL_{D_i}]$, resets its counter, and sends a control packet to $H_i, V_i,$ and D_i
Step5	$H_i, V_i,$ and D_i respectively resets $WL_{H_i}, WL_{V_i}, WL_{D_i}$ to zero and turn on/off adaptive antennas
Step6	Goto Step1

As shown in Figure 1, a local arbiter is connected to the wireless routers via metallic links. In this paper, we assumed a flit size of 64 bits with four flits in a packet. Hence, a packet takes 8 cycles to transmit through the wireless link. Therefore, each wireless router requires a 4-bit counter, and each arbiter requires a 7-bit counter. Since the size of the counters and the width of the metallic links are small, the overhead is insignificant; and thus, ignored in the performance evaluation.

3.4 Routing Mechanism of 256-core R-OWN

There are four clusters in a 256-core R-OWN where each cluster contains 16 routers and each router connects 4 cores. A core is represented by a 3-digit coordinate with *cluster*, *router*, and *core* number as follows: (cs, r, c) where cs is cluster, r is router, and c is core number. Thus, the total number of cores in R-OWN is $cs \times r \times c$, where $0 \leq cs \leq 3$, $0 \leq r \leq 15$, and $0 \leq c \leq 3$. Since cores communicate through routers, we drop the core index to identify a router.

Consider the R-OWN communication shown in Figure 2. For example, core $(0, 0, 0)$ and core $(0, 7, 2)$ both want to send packets to core $(1, 7, 3)$, and router $(0, 3)$ (H0) possess the adaptive wireless link of cluster 0. In other words,

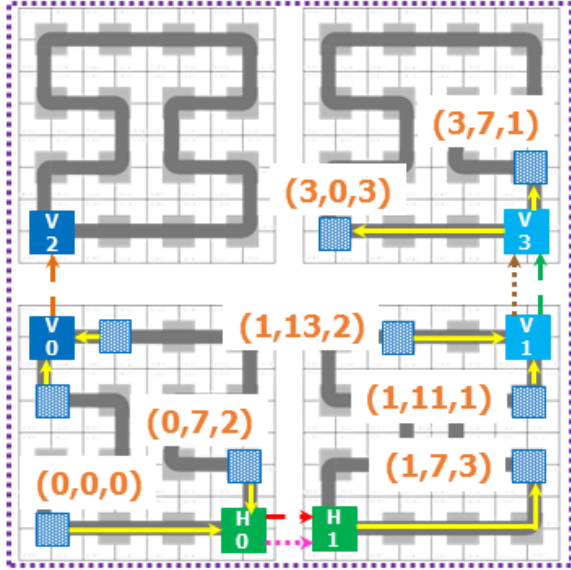


Figure 2: Communication mechanism of 256-core R-OWN. The large dotted line represents fixed wireless link, small dotted line represents adaptive wireless link and the solid line represents optical link. Routers of the same color talk to each other.

adaptive wireless link-F0 is connected to cluster 1 at this point of time. Both the cores will need to send a packet to the router H0 for inter-cluster wireless transmission. By modulating the wavelengths associated with router H0, one of the cores will send a packet first, and then the next core will send a packet. Assume both the packets are now sitting at the input buffer of router H0. Since two wireless links (one fixed, F01 and one adaptive, F0) are now connected to the wireless router H1 of cluster 1, these two packets will be sent concurrently using frequency channel F01 and F0. At the same reconfiguration time frame, for example, two cores of cluster 0 want to send packets to cluster 2 which requires the use of vertical wireless link (F02). Since only one wireless link is connected to cluster 2 from cluster 0, both the packets will contend for F02 at router V0 and packets will be transmitted serially. In contrast, say the adaptive wireless link of cluster 1 (F1) is pointing to cluster 3 as shown in Figure 2. Hence, core (1, 13, 2) and core (1, 11, 1) both will be able to send packets at the same time—using fixed wireless channel, F13 and adaptive wireless channel, F1—to their destination cluster 3 once the packets reach the wireless router V1. This is possible, because each cluster has its own adaptive wireless link which is configured based on the outgoing traffic from this cluster only. Now, consider core (1, 13, 2) send the packet first to destination core (3, 7, 1) using wireless link F13. Then if core (1, 11, 1) wants to talk to core (3, 0, 3), router V1 will use the wireless channel F1 instead of F13 as F13 was used last time. We chose to send packets using the adaptive and fixed wireless links alternatively to minimize contention. However, when a wireless router does not have access to the adaptive wireless link, we use the dedicated wireless link to communicate with the other clusters.

3.5 Deadlock Free Routing

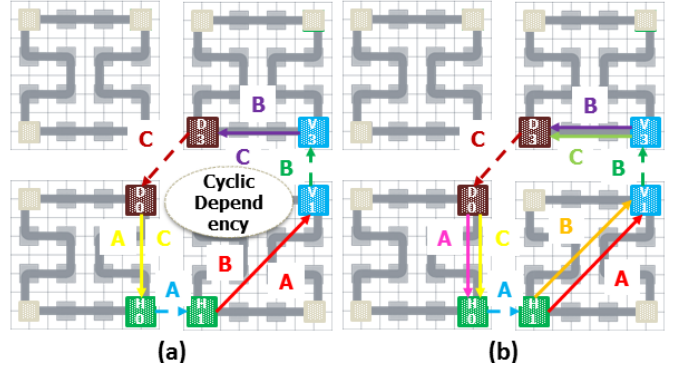


Figure 3: (a) Possible deadlock scenario in a 256-core R-OWN for simultaneous transmission of inter-cluster packets A, B, and C. (b) Proposed network with inclusion of new optical links to avoid deadlocks. A packet is marked with the color of the channel it is using.

Since R-OWN requires optical to wireless to optical domain transitions, cyclic dependency exists between the channels which may create deadlock. This is shown in Figure 3 (a) for three packets A, B, and C. Travel paths of packet A, B, and C are D0-H0-H1-V1, H1-V1-V3-D3, and V3-D3-D0-H0 respectively. Because packet A and C, or B and A, or C and B use the same optical link, deadlock may occur. To avoid this deadlocks, we have provided additional channels with usage restrictions. We assign new optical links for inter-cluster packets from the source router to the wireless router. However, on the destination cluster, packets use the optical links that were present before. As a result, for example, packet A and C take different optical links to travel from router D0 to router H0 which breaks the cyclic dependency. As shown in Figure 3 (b), the proposed network is deadlock-free which ensures all packet delivery. Since an optical waveguide can contain a maximum of 64 wavelengths and we can insert these additional optical links to the existing data waveguide, the tradeoff is increased power consumption.

4. PERFORMANCE EVALUATION

To evaluate the performance of the proposed architecture, we compare the 256-core R-OWN with CMesh [2], WCube [10], and OWN [15] architectures. We used Dsent v. 0.91 [16] to calculate the area and power of the wired links and routers for a bulk 45nm LVT technology. To simulate network performance for different types of synthetic traffic patterns such as uniform (UN), bit-reversal (BR), matrix transpose (MT), perfect shuffle (PS), and neighbor (NBR), we have used a cycle accurate simulator [8] keeping the router and core frequency same for all the networks. In order for a fair comparison between different topologies, we have kept the bisection bandwidth same for all the architectures by adding appropriate delay into the network.

4.1 Area Estimation

An architecture area consists of router and link (wired, optical, and wireless) area. As shown in Figure 4(a), R-OWN requires 12.47% less, 3.88% less, and 11.48% more area compared to WCube, CMesh, and OWN respectively. Since similar to OWN, R-OWN has lower number of input

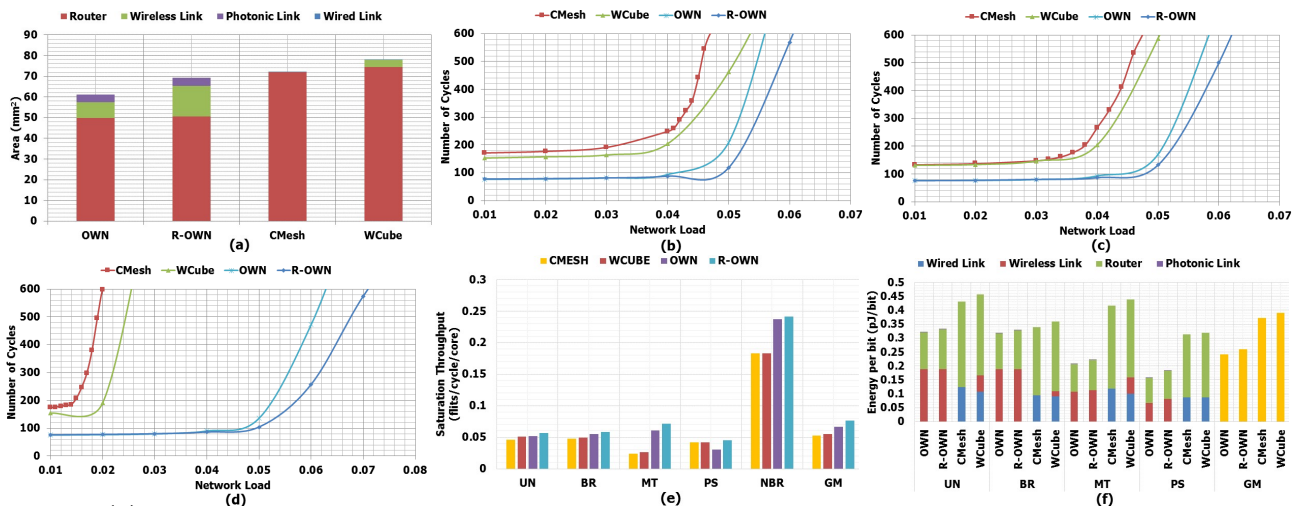


Figure 4: (a) Area comparison between the compared topologies that includes link and router area. Latency comparison between different networks are shown for (b) uniform traffic (UN), (c) bit-reversal traffic (BR), and (d) matrix transpose traffic (MT). (e) The saturation throughput of the comparing architectures with geometric mean (GM). (f) Energy per bit comparison between different topologies for various types of traffic. This energy calculation includes both leakage and dynamic components.

ports and the crossbar of the optical router is split into two, R-OWN requires less router area [15] compared to WCube and CMesh. However, R-own requires slightly higher router area than OWN due to a higher radix of the wireless router. For a wireless link, we have assumed the transmitter area as 0.42 mm^2 and receiver area as 0.20 mm^2 [9]. Because the number of wireless antenna in R-OWN is higher than both WCube and OWN, R-OWN requires highest wireless link area. In addition, since photonic link area consists of the power, data, and arbitration waveguide area, R-OWN requires more photonic link area compared to the wired link area of CMesh and WCube.

4.2 Saturation Throughput and Latency

In this section, we discuss the latency and saturation throughput of R-OWN compared to CMesh, WCube, and OWN. During simulation for WCube, to provide the best performance, we have optimized the threshold-distance to use the wireless link instead of wired link. In addition, we have kept the same number of VC and buffer per VC for all the architectures compared. Figure 4 (b, c, and d) shows the latency for the traffic types UN, BR, and MT respectively as a measure of the number of cycles in response to a varied network load. R-OWN performs the best compared to other architectures, because like OWN, R-OWN requires maximum of three hops to transmit to any part of the network whereas WCube and CMesh both have higher network diameter. Thus, for low network load, R-OWN and OWN both have similar latency since the contention in the network is low, and the improvement in R-OWN due to the reconfiguration is small. However, as the load increases, R-OWN performs better than OWN because it allocates the adaptive wireless channels efficiently to the routers that experience more traffic.

Figure 4 (e) shows the saturation throughput for traffic types UN, BR, MT, PS, and NBR where GM is the geometric mean. Since OWN and R-OWN both have the lowest diameter, they have the highest saturation throughput for UN, BR, and MT. However, since R-OWN adapts with the

network load pattern, it has the higher throughput than OWN. For PS, utilization of wireless links are diverse. This causes the saturation throughput of OWN to fall as certain wireless link is over utilized where the others are under-utilized. Thus, the improvement of R-OWN with respect to OWN is the highest for PS. Therefore, R-OWN has a saturation throughput which is 13.06%, 27.48%, and 31.08% higher than OWN, WCube, and CMesh respectively.

4.3 Energy

To calculate the wired link energy consumption, we have multiplied the number of wired link traversals which were collected from the cycle accurate simulation to the corresponding wired link energy calculated using Dsent [16]. For fairness, we have assumed a fixed 1 pJ/bit energy consumption for a wireless link in WCube, OWN, and R-OWN. To calculate the optical link energy consumption for OWN and R-OWN, we have considered the power, data, and arbitration waveguide energy consumption. We have calculated the router energy consumption per flit and taken into account the crossbar splitting of the optical router.

Figure 4 (f) shows the energy per bit comparison for UN, BR, MT, and PS traffic patterns with the geometric mean. Since the number of wireless link traversal for both WCube and OWN is less compared to R-OWN, R-OWN consumes more wireless link energy. The difference in energy consumption of OWN and R-OWN is the highest for PS traffic, because for PS traffic the adaptive wireless links are well utilized which is also reflected in their saturation throughput (Figure 4(e)). R-OWN consumes lower router energy than CMesh and WCube due to the lower hop requirement, lower number of input ports with higher number of output ports, and splitting of the crossbar. However, R-OWN requires higher router energy compared to OWN due to the increase of wireless router radix.

From the Figure 4(f), we can see that more than 50% of the consumed energy for OWN and R-OWN is wireless link energy. This is because OWN and R-OWN both use wireless for all the inter-cluster transmission whether they

are neighbor or not, and the wireless link energy per bit requirement is comparatively high. Although, WCube consumes lower wireless link energy than R-OWN, this is offset by higher wired link and router energy consumption. The end result is that R-OWN consumes 43.70% less, 50.56% less, and 6.95% more energy/bit than CMesh, WCube, and OWN respectively.

5. CONCLUSIONS

The proposed architecture R-OWN combines the benefits of photonic and wireless interconnect technologies to scale future CMPs. Our results show that R-OWN requires 12.47% less, 3.88% less, and 11.48% more area compared to WCube, CMesh, and OWN respectively. With advances in CMOS technology, the transceiver and photonic link area may reduce which will benefit R-OWN. In addition, R-OWN consumes 43.70% less, 50.56% less, and 6.95% more energy/bit than CMesh, WCube, and OWN respectively. Furthermore, R-OWN shows lowest latency than all the comparable topologies for the simulated traffic patterns. R-OWN has 13.06%, 27.48%, and 31.08% higher saturation throughput compared to OWN, WCube, and CMesh respectively. Moreover, R-OWN leverages the benefits of two technologies while balancing their drawbacks and utilize the flexibility of creating dynamic wireless connections.

6. ACKNOWLEDGMENT

This research was supported by the grant funding numbers: CCF-1054339 (CAREER), CCF-1420718, CCF-1318981, ECCS-1342657, CCF-1513606, CCF-1547034, CCF-1547035, CCF-1565273, and CCF-1600820. We would like to thank the anonymous reviewers for their valuable feedback.

7. REFERENCES

- [1] MPPA Manycore. http://www.kalrayinc.com/IMG/pdf/FLYER_MPPA_MANYCORE.pdf. Accessed: 06-12-2016.
- [2] J. Balfour and W. J. Dally. Design tradeoffs for tiled cmp on-chip networks. In *Proceedings of the 20th annual international conference on Supercomputing*, pages 187–198. ACM, 2006.
- [3] J. Cunningham, R. Ho, X. Zheng, J. Lexau, H. Thacker, J. Yao, Y. Luo, G. Li, I. Shubin, F. Liu, et al. Overview of short-reach optical interconnects: from vcsels to silicon nanophotonics.
- [4] D. DiTomaso, A. Kodi, S. Kaya, and D. Matolak. iwise: Inter-router wireless scalable express channels for network-on-chips (nocs) architecture. In *High Performance Interconnects (HOTI), IEEE 19th Annual Symposium on*, pages 11–18. IEEE, 2011.
- [5] D. DiTomaso, A. Kodi, D. Matolak, S. Kaya, S. Laha, and W. Rayess. A-winoc: Adaptive wireless network-on-chip architecture for chip multiprocessors. *Parallel and Distributed Systems, IEEE Transactions on*, 26(12):3289–3302, Dec 2015.
- [6] A. Ganguly, K. Chang, S. Deb, P. P. Pande, B. Belzer, and C. Teuscher. Scalable hybrid wireless network-on-chip architectures for multicore systems. *Computers, IEEE Transactions on*, 60(10):1485–1502, 2011.
- [7] A. Joshi, C. Batten, Y. J. Kwon, S. Beamer, I. Shamim, K. Asanovic, and V. Stojanovic. Silicon-photonic clos networks for global on-chip communication. In *Proceedings of the 3rd ACM/IEEE International Symposium on Networks-on-Chip*, pages 124–133. IEEE Computer Society, 2009.
- [8] A. Kodi and A. Louri. A system simulation methodology of optical interconnects for high-performance computing systems. *J. Opt. Netw.*, 6(12):1282–1300, 2007.
- [9] A. K. Kodi, M. A. I. Sikder, D. DiTomaso, S. Kaya, S. Laha, D. Matolak, and W. Rayess. Kilo-core wireless network-on-chips (nocs) architectures. In *Proceedings of the Second Annual International Conference on Nanoscale Computing and Communication, NANOCOM' 15*, pages 33:1–33:6. ACM, 2015.
- [10] S. B. Lee, S. W. Tam, I. Pefkianakis, S. Lu, M. F. Chang, C. Guo, G. Reinman, C. Peng, M. Naik, L. Zhang, et al. A scalable micro wireless interconnect structure for cmps. In *Proceedings of the 15th annual international conference on Mobile computing and networking*, pages 217–228. ACM, 2009.
- [11] D. Matolak, A. Kodi, S. Kaya, D. DiTomaso, S. Laha, and W. Rayess. Wireless networks-on-chips: architecture, wireless channel, and devices. *Wireless Communications, IEEE*, 19(5):58–65, Oct 2012.
- [12] R. Morris, A. Kodi, and A. Louri. Dynamic reconfiguration of 3d photonic networks-on-chip for maximizing performance and improving fault tolerance. In *Microarchitecture (MICRO), 45th Annual IEEE/ACM International Symposium on*, pages 282–293, Dec 2012.
- [13] A. Olofsson, R. Trogan, O. Raikhman, and L. Adapteva. A 1024-core 70 gflop/w floating point manycore microprocessor. In *Poster on 15th Workshop on High Performance Embedded Computing HPEC*, 2011.
- [14] 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*, volume 37, pages 429–440. ACM, 2009.
- [15] M. I. Sikder, A. K. Kodi, M. Kennedy, S. Kaya, and A. Louri. Own: Optical and wireless network-on-chip for kilo-core architectures. In *High-Performance Interconnects (HOTI), IEEE 23rd Annual Symposium on*, pages 44–51, Aug 2015.
- [16] C. Sun, C. H. Chen, G. Kurian, L. Wei, J. Miller, A. Agarwal, L. S. Peh, and V. Stojanovic. Dsent-a tool connecting emerging photonics with electronics for opto-electronic networks-on-chip modeling. In *Networks on Chip (NoCS), Sixth IEEE/ACM International Symposium on*, pages 201–210. IEEE, 2012.
- [17] 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*, volume 36, pages 153–164. IEEE Computer Society, 2008.