

LumiNoC –A Novel Design for a Power-Efficient, Performance Oriented Photonic Network-On-Chip

Divya Bhadrakumar¹, Anoop T.R², Seema Padmarajan³

¹(Electronics and Communication Department, Sree Narayana Gurukulam College of Engineering, India)

²(Electronics and Communication Department, Sree Narayana Gurukulam College of Engineering, India)

³(Electronics and Communication Department, Sree Narayana Gurukulam College of Engineering, India)

Abstract : Chip Multiprocessors (CMPs) fabricated with the provision for numerous, comparatively simpler processor cores, overcome the limitations in parallelism, clock speed and design cost encountered in large uniprocessors. This paper explores a paradigm shift in computer architecture to incorporate CMPs along with Network-on-Chips (NoCs); the latter addresses the many challenges involved in interconnecting the increased number of cores and scaling. The challenges posed by the traditional NoCs in terms of throughput, design, topology, power dissipation, bandwidth and latency optimization and on-chip temperature are eliminated through the use of Photonic NoCs (PNoCs). Silicon nanophotonics comes to the rescue by providing a replacement for electronic on-chip interconnects with its high bandwidth and better latency records. In this paper, a new architecture for NoC, referred to as LumiNoC is discussed, that is optimized for better performance and efficient power management, by conceiving partitions of the network into subnets, thereby increasing the efficiency. The simulations for router and channel architectures have been carried out in Xilinx ISE Suite 14.2.

Keywords - CMP, LumiNoC, nanophotonics, NoC, PNoC

I. INTRODUCTION

With increasing application complexity and further innovations and improvements in process technology, it lead to the development of the ‘Chip Multiprocessors’ or CMPs. This allowed not just a single core, but multiple cores (10s to 100s) to be implemented on a single chip. The studies in this area were triggered early, in 1967, with the proposal of the Amdahl’s law or argument. Amdahl’s law was used to predict the theoretical maximum speed-up using multiple processors. He said that the speed-up of a program using multiple processors in parallel computing is limited by the time needed for the sequential fraction of the program to be executed. Network-on-Chip (NoC) has been an emerging paradigm in Very Large Scale Integrated (VLSI) systems that are implemented on a particular silicon chip. A typical NoC construct has multiple point-to-point data links interconnected by switches (or routers), such that messages can be relayed from any source module to any destination module over several links, by making routing decisions at switches, perhaps by some routing algorithm. NoCs are being adopted due to its inherent advantages of scalability, flexibility, modularity and increased productivity. It offers an extremely well-controlled structure that provides a partition between computation and communication fields of a system.

Interconnects are used to connect components on a VLSI chip, chips in multichip modules, and multichip modules on a system board. Traditionally used electrical interconnects, however, are severely limited by power, bandwidth and latency constraints. These constraints place practical limits on the viability of future CMP scaling. So, the design of NoC architecture introduces challenges in terms of offered throughput, layout, power efficiency and topology. Through various studies, statistics show that power dissipation of NoC is more than 25% of the overall power. Consequently, the limited on-chip power budget will have to be carefully distributed between computation and communication activities, by allowing for higher power devoted for cores and lower power dissipated by the NoC. This improves the system performance. Besides, metallic or electrical interconnects have a number of issues such as crosstalk, propagation delay, parasitic capacitance and inductance and electromigration. So another viable alternative has been proposed – optical interconnects; they are reliable, noise-free, provide high speed transmission capabilities and larger bandwidths. Such a NoC that uses optical interconnects is called as a photonic NoC (PNoC).

Processing directly in the optical domain has several limitations, primarily due to the lack of optical memories or equivalent optical RAM. This calls for combining photonic communication with electronic control, or a „hybrid approach“. The photonic interconnection network is used to transfer large data messages between the cores whereas the electronic interconnection network is packet switched and is used to carry small control packets. In this paper, a new architecture for NoC, referred to as LumiNoC is discussed, which is optimized for

better performance and efficient power management, by conceiving partitions of the network into subnets, thereby increasing the efficiency. The design incorporates a channel sharing arrangement that utilizes waveguides and wavelengths to arbitrate data transmission.

LumiNoC effectively addresses resource overhead due to over provisioning and power consumption, and avails an intensive data transmission at higher bandwidth with much reduced latencies. LumiNoC offers three main contributions: a channel sharing arrangement; a purely photonic, distributed, dynamic channel scheduling; leveraging the same wavelengths for channel arbitration and parallel data transmission.

The paper has been organized into seven sections in the following manner. Section II elaborates on the literature survey. It describes in detail the concept of PNoC. Section III illustrates the proposed methodology employed by the LumiNoC architecture in detail. Section IV portrays the simulation result. Section V condenses the advantages and disadvantages of the proposed methodology. It also mentions the possible future scope in this area of study.

II. LITERATURE SURVEY

Further scaling of the Moore's Law has been made possible with multi-core processors (or CMPs, fabricating chips with multiple cores, allowing for efficient parallel processing), which have helped overcome the bottlenecks of heat dissipation and data synchronization. Predictions are that by 2020, the number of cores will increase by multiple hundred cores per chip. But, however, the performance will again be limited by the Amdahl's Law, and the maximum number of cores is expected to be around a thousand cores [3].

There are three primary metrics [3] which govern the interconnect technology – bandwidth, latency and energy efficiency. The insertion of photonics in the on-chip interconnect structure can be potentially influenced by the unique advantages of optical communication. It furthermore improves the capacity, bit transparency, and lowers energy consumption. A network-on-chip incorporating optical links as interconnects are termed as Photonic Network-on-Chip or PNoC. PNoC is an effective solution to reduce the overall power budget. Optical waveguides offer lower loss and bit-rate transparency. With this combined, a PNoC can potentially deliver considerably higher bandwidth and lower latencies with significantly lower power dissipation as compared to a traditional interconnection network based only on electronic signalling.

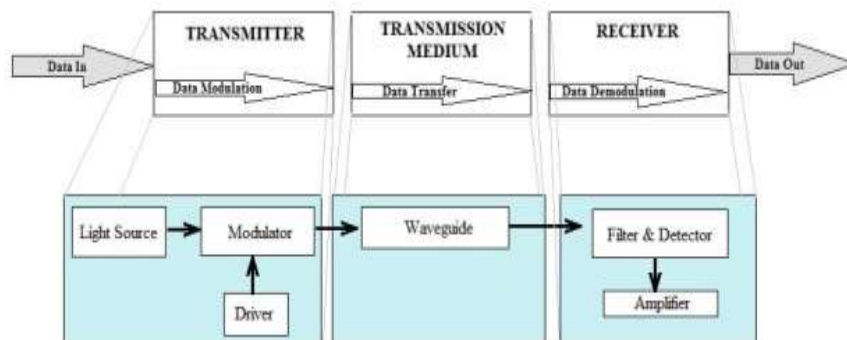


Fig. 1: High-level photonic communication link with photonic schematic enlargements [4]

Fig. 1 shows a high-level illustration of a generic photonic link [4], with schematic enlargements that identifies the key constituent components of the system. In order to transmit the data (stored in electrical form at sender) to the receiver, the data must be modulated onto the light. The modulated light is then coupled to a transmission medium, which transports the optical signals to the receiver. The receiver consists of a filter (used for separating the individual communication channels) and detectors (for converting the light back to electrical signal). The advantages of photonic interconnect over the electrical interconnects are (1) better energy efficiency and lower link latency, (2) bit rate transparency, (3) superior raw bandwidth and bandwidth density, (4) high carrier frequency of optical signals allows them to be guided by extremely low-loss dielectric waveguides, (5) photonic signals do not suffer from propagation loss and distortion due to resistance of the metal waveguides, (6) they are relatively immune to reflection and crosstalk, (7) all these factors combine to give photonic links the ability to communicate over longer physical distances and at significantly higher data rates without major latency or power penalties, (8) given its ability to efficiently communicate over longer distances,

photonics allows designers to use flatter networks, with uniform communication latency without compromising on scalability.

Each core in the CMP is composed of a network interface and a gateway. The gateway performs electronic/optical and optical/electronic (E/O and O/E) conversions, to communicate with the control network and perform clock synchronization and recovery, as well as serialization and deserialization of messages. Communication between any two nodes in a PNoC is set-up through the network in three main steps [7], [9]: path-setup process, photonic transmission process and finally, the path-tear down process. The path set-up packet (electronic control packet) is routed on the electronic network, acquiring and setting-up a photonic path for the message. The photonic messages are then transmitted without buffering once the path has been acquired. This approach is similar to an optical circuit switching.

Consider a processor at node A transmitting data to a memory at node B. A path-setup packet is sent on electronic control network. It includes information on the destination address of node B and additional control information such as priority, "1-hot source" address and flow id. The path set-up packet is routed in the electronic control network. It reserves the photonic switches as it travels along the path. The next hop is determined at every router in the path according to some routing algorithm used. By the time the path-setup packet reaches the destination, the photonic path is reserved. Then, a fast light pulse is sent on the photonic path from node B to node A to indicate that the path has been reserved. The photonic message is now sent from node A, following the path from switch to switch, until it reached node B. Once the packet reaches node B, the message transmission is completed. Next, the path-teardown packet is sent from node B to node A on the electronic control network to relieve the path. Also, the photonic message is checked for errors and a small acknowledgement packet is sent from node B to node A on the electronic control network.




Fig. 2: Photonic switching element: (a) ON state: a passive waveguide; (b) OFF state: light is coupled into rings and forced to turn

The fundamental building block of the photonic network is the photonic switching element (PSE) [7], [2]. The PSE is based on a microring resonator structure. The PSE is a waveguide intersection, positioned between two ring resonators (Fig. 2). The PSE demonstrates a switching action. The switch is turned ON by injecting an electrical current into the p-n contacts surrounding the rings. On doing so, the resonance shifts, such that the incident light signals are off resonance, and the light passes through the waveguide intersection uninterrupted (Fig. 2 (a)). In the OFF state, the resonance frequency of the rings coincides with the wavelength on which the optical data stream is modulated. The light is coupled into the rings and onto the perpendicular waveguide, making right-angle turns (Fig. 2(b)). The PSEs are interconnected by silicon waveguides which carry photonic signals. The PSEs are organized as groups of four. An electronic router (ER, an electronic circuit) controls the quadruple, forming a 4x4 switch. The electronic router is connected to the network with metal lines. The 4x4 switches are interconnected by the inter-PSE waveguides.

During the path-setup step, the electronic router receives the control packets. The ER processes them and sends them to their next hop, while switching the PSEs ON and OFF appropriately. After the path-setup process sets up the photonic path through a sequence of ERs, the data transmission step is carried out wherein the data message is routed over a chain of PSEs. The previously switch is internally blocking, which requires specific routing algorithms to improve performance. A non-blocking switch alleviates this problem by increasing the number of internal paths within the switch.

III. PROPOSED WORK

A novel, hybrid architecture for PNoC is discussed in this paper, that combines a broadband photonic circuit-switched network with an electronic overlay packet-switched control network [6]. It presents a flexible alternative for data transmission on-chip, though larger messages being communicated over the photonic

network and the shorter control messages being delivered electronically, with minimal power consumption and greater efficiency. Fig. 3 shows a minor CMP consisting of three compute tiles, connected using PNoC. Each one of the tiles or nodes comprises of its own processor core, private caches, and a router connecting it to the photonic network [1]. The photonic channel connects all the three nodes together in the network. The photonic channel consists of microring resonators (MRR), photodetectors (PD, denoted as small circles) and silicon waveguides (represented as black lines connecting the circles). Transceivers are also included (represented as small triangles); they mark the periphery between the electrical domain and photonic domain.




Fig. 3: A three-node fully connected photonic crossbar

1. Basic Working Concept

If a sending node and the receiving node do not dwell on the same subnet, then the transmission requires a hop through an intermediate node's router (or switch), in which event the transmission suffers a comparatively longer delay due to the latency offered by the router and also the delay on account of the E/O and O/E conversions at the network gateways. So, in order to eliminate the overheads of the photonic waveguide crossings for the given orthogonal set of horizontal and vertical subnets, the waveguides are deposited into two layers with orthogonal routing. The LumiNoC architecture is improved over the conventional crossbar architecture by incorporating multiple tiles on the same photonic channel though sharing. However, by allocating many wavelengths on to the same waveguide using multiplexing schemes like WDM leads to high waveguide losses. To overcome this problem, in the suggested LumiNoC architecture, the number of wavelengths has been limited to a few frequencies per waveguide and waveguide count per subnet is also increased.

Each node or tile consists of a transmitting ring ("TX") as well as a receiving ring ("RX"). The number of transmitting and receiving rings is equal to the total number of wavelengths multiplexed in the waveguide, represented as " \wedge ". The optical signal propagates unidirectional along the waveguide starting from the source at off-chip laser. So, as can be seen in Fig. 4, the transmitting ring ("TX") of each node is connected in series to the data-send path (denoted by the straight blue line), after which the receiving ring ("RX") of each node is connected in a similar manner to the data-receive path (denoted by the dotted red line). So as can be seen, the modulation by any node can be received by any other node in the subnet, hence this arrangement is called as a "double-back" waveguide layout. Moreover, the transmitting node can also receive its own modulated signal, a feature leveraged by a collision detection scheme at the arbitration and data transmission phases.

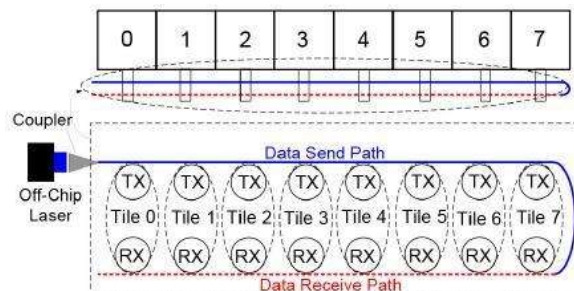


Fig. 4: One-row subnet of eight nodes. Circles (TX and RX) denote groups of rings; one dotted oval represent a tile

At any given time during data transmission, only a single transmitting node can modulate on all the wavelengths, and only a single receiving node can be tuned to all of the wavelengths. On the other hand, during arbitration phase, the receiving rings of all the nodes in the subnet will be tuned to a particular non-overlapping set of wavelengths. At any given time, a given multi-wavelength channel comprising of “N” nodes can be in any one of the following three states: idle state, arbitration phase, data transmission phase. In the idle state, the network is in a quiescent state, with all wavelengths unmodulated. The arbitration phase is said to occur if more than one node is simultaneously trying to acquire the control over the channel, by modulating copies of the arbitration flag. Once a particular sender has acquired control over the channel the data transmission phase is said to occur. In this state, the sender modulates all channel wavelengths in parallel with the data to be transmitted. The arbitration flag consists of the following three fields:

- The destination node address (D0 – D1)
- A bimodal packet size indicator (Ln)
- The “1-hot” source address (S0 – S2)

In the architecture implemented, the network consists of three nodes. Two bits (D0-D1) are sufficient to represent the address of each node. The bimodal packet size indicator Ln, indicates the length (number of bits) of the data transmitted. For each of the nodes in the network, 1-bit is allocated in the arbitration flag to indicate the nodes transmitting data at a particular given time. Logic high in the “1-hot” source field indicates that the corresponding node is trying to send data or access the channel. The time duration of the arbitration flag is represented as t_{arb} .

Before sending a packet of data, the transmitting node ensures that all ongoing processes are completed. Then, it modulates a copy of its arbitration flags onto an appropriate wavelength for each of the nodes in the network. When multiple transmitting nodes send overlapping arbitration flags, the “1-hot” precondition is violated and all the nodes in the network will detect a collision. Due to the “double-back” arrangement of the architecture, a provision for self-reception of the arbitration flag by the sending node is availed. However, if the arbitration flag is corrupted, a conflict is perceived and any data that may already have been sent over the channel is ignored and all the transmitting nodes enter the dynamic channel allocation regime.

On detecting an arbitration flag, the receiving node performs the following steps. If arbitration flag is uncorrupted, and the incoming data is destined for the corresponding receiving node, all the RX rings are enabled to capture the data. If the arbitration flag is uncorrupted but the incoming data is not destined for the corresponding receiving node, all the RX rings are detuned for the time duration of the data to allow the recipient the sole access. If however, collision is detected, that is, the arbitration flag is corrupted, the receiver circuit enters the dynamic channel allocation phase.

When a corrupted arbitration flag is detected, a collision is detected and we enter the dynamic channel scheduling phase. When a conflict is intuited, all the nodes identify the conflicting transmitting nodes. After identification a fair, dynamic scheduling of the conflicting nodes is done as per some priority (according to transmitting node index). As per the now allocated schedule, the transmitting nodes transmit their data one after another, after transmitting an arbitration flag over the channel. All nodes tune in to receive the data. The channel is occupied until the last sender transmits its data packet. After the dynamic channel allocation is completed, the network enters the idle state or quiescent state. Next, any node can attempt to transmit arbitration flag to acquire channel to transmit its data. The data transmission phase succeeds the arbitration phase, after the channel is allocated to a particular pair of nodes in the network. During this phase, the transmitting node sends its data over the photonic channel to its corresponding destination or receiving node in the network. During the transmission phase all the wavelengths allocated in the waveguides is utilized for parallel data transmission, to ensure higher throughput and greater efficiency. The credit return phase is used to provide acknowledgement of the success of the data transmission phase to the transmitting node.

2. Router Architecture

Every router in the LumiNoC architecture serves as both an entry point as well as an intermediate node for data transmission. If a node or tile intends to send a data packet to another node or tile on the same vertical or horizontal subnet of the network, then the router acts as a switch, causing the packet to be switched from the electrical input port to the vertical photonic output port after carrying out E/O conversion. On the other hand, if

the data is to be sent from one node to another node on another subnet of the network, then the router acts as an intermediate node; the packet is first routed through the intermediate node via the horizontal subnet before being routed onto the vertical subnet. The Fig. 5 shows the schematic of the router module.

The top level router module consists of the following six sub blocks: three FIFO modules, a synchronizer module, a register module and a finite state machine (FSM). The FIFO works on the system clock and is reset with a synchronous active low reset. The synchronizer module provides the synchronization between the router FSM and the router FIFO modules. It provides faithful communication between the single input port and three output ports. The register module implements four internal registers in order to hold the header byte, the FIFO full state byte, the internal parity and the packet parity byte. The FSM module is the controller circuit of the router. This module generates all the control signals when a new packet is received by the router. These control signals are used by other design components in order to transfer the packet to the output port.

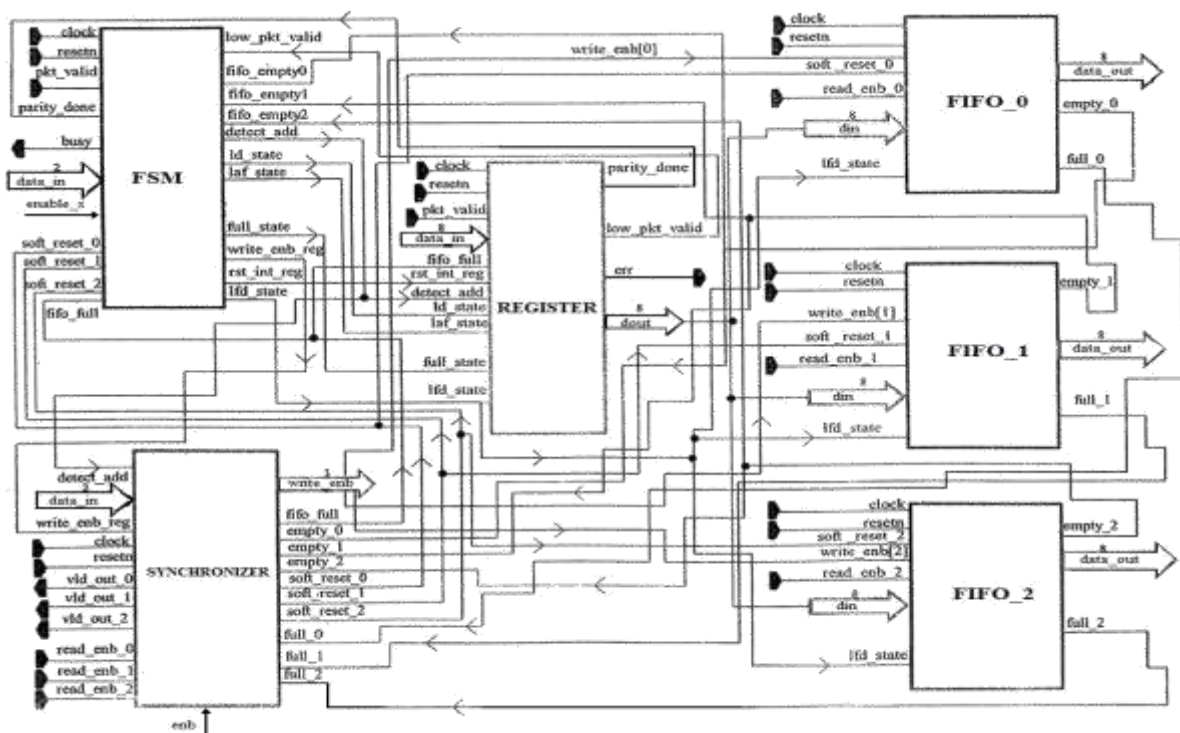


Fig. 5: Router top-level RTL schematic




Fig. 6: FSM State Diagram

3. Channel Architecture

The channel architecture consists of two different blocks-one for normal mode and another for dynamic channel allocation (DCA) mode. When only one source is requesting for the channel (i.e. in keeping with the “1-hot” source address), then the normal mode is enabled, while disabling the DCA mode. However, when more than one source is trying to access the channel control, the DCA mode is enabled while disabling the normal mode. In the DCA mode, the channel is allocated one at a time to each of the requesting sources as per some priority assigned. After all the sources complete their transmission, then the system enters the idle mode where no operation is being carried out – a quiescent state. While one node is working, a busy signal is given to the other nodes, indicating that the channel is occupied and these nodes remain idle until the busy signal becomes low. A signal for indicating whether the data is being transmitted between the nodes in the network is also provided.

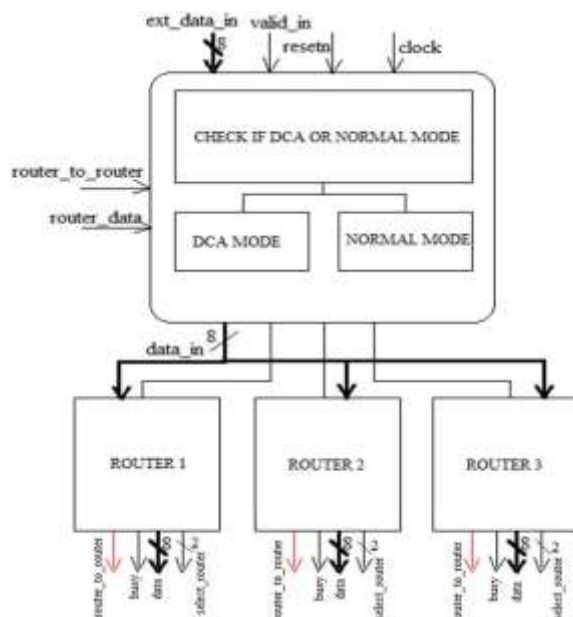


Fig. 7: Channel Architecture

IV. SIMULATION RESULTS

The simulation has been carried out in Xilinx ISE Design Suite (Version 14.2) platform, programmed in Verilog Hardware Description Language (HDL). The results have been obtained for the router section and the channel section as illustrated in Fig. 8 and Fig. 9, respectively. The output has been obtained for the three tile architecture discussed in section III.



Fig. 8: Simulation result for router module




Fig. 9: Simulation result for channel module

V. CONCLUSION & FUTURE SCOPE

In conclusion, the LumiNoC architecture addresses the issues of traditional NoCs by adopting a shared-channel and in-band arbitration mechanism. It efficiently utilizes power, achieving a high performance, and scalable interconnect with extremely low latency. LumiNoC offers lower latency at low loads and higher throughput as when compared to other PNoC architectures. On the other hand, there are some limitations to the architecture as well. This is primarily due to the use of off-chip laser. Although the off-chip laser does not impact the processor power budget, it does impact the system energy efficiency. Off-chip lasers also add-on additional coupling losses and packaging expenses. However, off-chip lasers provide for easy replacement and temperature stability. It has also been discovered through studies that optical interconnects are not feasible with memory bandwidth also included.

It can be deduced that with the ITRS roadmap as the foundation, electrical interconnects will not be capable of sustaining itself in the VLSI chips, with further scaling of on-chip components and interconnects. As a result within the next decade electrical interconnects will fade in to insignificance, and optics will emerge as a better choice. Furthermore, optics is widening its scope with novel ideas in the area of interconnects, for

instance, in carbon nanotube based interconnects. So, there is high potential in the research being carried out in the field of PNoC. Additionally, for on-chip global interconnects, such as those used between multiple cores, optics seems to be the near-term interconnect technology of choice. Moreover, different topologies and switching methods to further improve the parameters can also be explored.

Acknowledgements

My profound and sincere expression of thanks and gratitude to my mentors Mr. Anoop T.R, Ms. Seema Padmarajan, Ms. Malini Somanand and Ms. Soniya Peter for their unstinted support, help and guidance. My sincere thanks to one and all who may not find a mention but have always wished only success.

REFERENCES

Transaction Papers:

- [1] Cheng Li, Mark Browning, Paul V. Gratz, and Samuel Palermo, LumiNOC: A Power-Efficient, High-Performance, Photonic Network-on-Chip, *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, Vol. 33, No. 6, June 2014.
- [2] A. Shacham, K. Bergman, and L. P. Carloni, Photonic networks-on-chip for future generations of chip multiprocessors, *IEEE Trans. Comput.*, vol. 57, no. 9, pp. 1246–1260, Sep. 2008.

Journal Papers:

- [3] Martijn J. R. Heck, and John E. Bowers, Energy Efficient and Energy Proportional Optical Interconnects for Multi-Core Processors: Driving the Need for On-Chip Sources, *IEEE Journal of Selected Topics in Quantum Electronics*, Vol. 20, No. 4, July/August 2014.

Books:

- [4] Christopher J. Nitta, Matthew Farrens, Venkatesh Akella., *On-Chip Photonic Interconnects: A Computer Architect's Perspective* (Morgan & Claypool Publishers, October 2013).

Proceedings Papers:

- [5] Assaf Shacham, and Keren Bergman, Building Ultralow-Latency Interconnection Networks Using Photonic Integration, published by the *IEEE Computer Society*, July/August 2007.
- [6] A. Krishnamoorthy et al., Computer systems based on silicon photonic interconnects, *Proc. IEEE*, vol. 97, no. 7, pp. 1337–1361, Jul. 2009.
- [7] A. Shacham, K. Bergman, and L. P. Carloni, Photonic NoC for DMA communications in chip multiprocessors, in *Proc. 15th Annu. IEEE HOTI*, Stanford, CA, USA, 2007, pp. 29–38.
- [8] Wim Bogaerts, Peter De Heyn, Thomas Van Vaerenbergh, Katrien DeVos, Shankar Kumar Selvaraja, Tom Claes, Pieter Dumon, Peter Bienstman, Dries Van Thourhout, and Roel Baets, Silicon microring resonators, *Laser Photonics Rev.* 6, No. 1, 47–73 (2012).
- [9] Michele Petracca, Benjamin G.Lee, Keren Bergman and Luca P.C., Photonic NoC: System level design exploration, for *IEEE Xplore Computer Society*, July/August 2009.