Building Large Scale Networks using Power Efficient Nanophotonic Crossbars

Randy W. Morris, Jr. and Avinash K. Kodi
Electrical Engineering and Computer Science, Ohio University, Athens, OH 45701
E-mail: rmorris@cs.ohiou.edu,kodi@ohio.edu

Abstract: We propose two power efficient nanophotonic crossbar implementations for a 256-core nanophotonic NoC called E-PROPEL. Our results indicate E-PROPEL consumes 3× less power and increases performance by 2× when compared to leading electrical networks.

OCIS codes: (200.4650) Optical Interconnects; (250.6715) Switching

1. Introduction

High performance and low power network-on-chips (NoCs) will be required for future chip multi-processors (CMPs) which will be comprised of 10's, 100's, and even 1000's of cores. It is well-known that metallic NoCs will have difficulty in providing the needed bandwidth due to crosstalk, impedance mismatch, electromagnetic interference and power dissipation [1,2]. One potential solution is to use nanophotonic technology to overcome the limitation of metallic interconnects. With the recent surge in complementary-metal-oxide semiconductor (CMOS) compatible devices with small footprints (~10μm) and low power dissipation (~100 μW), nanophotonics are now a potential solution for on-chip interconnects [1,2].

In this paper, we propose a 64-core nanophotonic NoC called PROPEL and two 256-core versions called E-PROPEL (extended-PROPEL) and RE-PROPEL (reduced-extended-PROPEL). PROPEL uses optical interconnects for long distance inter-router communication and electrical switching within the routers. This reduces the power dissipation on long inter-router links while electrical switching provides flow control to prevent buffer overflow. In addition, we also propose two nanophotonic crossbars used to construct E-PROPEL and RE-PROPEL. In E-PROPEL, we allow for full connectivity between each PROPEL tiles and in RE-PROPEL, we restrict the connectivity between each PROPEL network. We compare E-PROPEL and RE-PROPEL to the mesh topology using synthetic traffic. Our results indicate the following: (1) E-PROPEL dissipates 3× less power than the mesh topology, (2) E-PROPEL increases performance by about a factor of 2 when compared to mesh, and (3) the nanophotonic crossbar designs are both area (0.0567 mm^2) and power (-4.32 dB loss) efficient.

2. PROPEL

Figure 1(a) shows the layout of PROPEL [3]. PROPEL consists of 16-tiles, each consisting of 4 cores, connected in a grid fashion with 4-tiles in x- and y-directions. Optical interconnects are used in two dimensions along the grid similar to an electronic 2D mesh or torus. Figure 1(b) shows the top x-direction tiles 0 to 3. Every tile modulates the same wavelength into a different waveguide. Each destination tile is associated with a waveguide called as the *home channel*. For example, tile T(0,0) has four modulators (ring resonators), all of which are resonant with the wavelength λ_0. Three λ_0 transmissions from tile T(0,0) are used to communicate with the other three tiles T(1,0), T(2,0) and T(3,0) on their home channel waveguides. The fourth resonant wavelength will be used to communicate with the memory bank. The wavelength assignment is repeated for tiles T(1,0), T(2,0) and T(3,0) .

3. E-PROPEL Implementation

The proposed E-PROPEL, which can be scaled up to 256 cores is shown in Figure 2. E-PROPEL is designed by replicating four PROPELs and connecting corresponding tiles on different PROPELs through a crossbar. In the figure, each crossbar is numbered to show it corresponding tiles it connects. For example, crossbar (0,0) connects all tiles T(0,0) together and crossbar (3,0) connects all the tiles T(3,0) together. In RE-PROPEL, we restrict inter-cluster communication to the top and bottom tiles of each cluster to reduce area requirements. E-PROPEL is designed as a fatree with multiple roots to provide scalable inter-cluster bandwidth. E-PROPEL requires 16 crossbars and RE-PROPEL requires 8 crossbars for connecting the tiles from different clusters.

4. Crossbar Implementation

A nanophotonic crossbar allows for N×N simulation communications to take place without requiring arbitration, which has great advantages for time sensitive network design such as E-PROPEL. Figure 3(a) shows the crossbar functionality of a 64-wavelength 4×4 crossbar which are used in the construction of E-PROPEL. Here, the wavelengths are indicated as $\lambda_{(a\text{-}b)}^{(c)}$, where *a-b* indicate the wavelength range and *c* indicates the input port. In the figure, consider input port 0. All input wavelengths are indicated as $\lambda_{(0-15)}^{(0)}$, $\lambda_{(16-31)}^{(0)}$, $\lambda_{(32-47)}^{(0)}$ and $\lambda_{(48-63)}^{(0)}$. After traversing the series of ring resonators, the wavelengths $\lambda_{(0-15)}^{(0)}$ arrive at output port 0, $\lambda_{(16-31)}^{(0)}$ at output port 1, $\lambda_{(32-47)}^{(0)}$ at output port 2 and $\lambda_{(48-63)}^{(0)}$ at output port 3. This enables a 1 × N switching functionality per waveguide. Now consider input port 1. All input wavelengths are indicated as $\lambda_{(0-15)}^{(1)}$, $\lambda_{(16-31)}^{(1)}$, $\lambda_{(32-47)}^{(1)}$ and $\lambda_{(48-63)}^{(1)}$. After traversing the series of ring resonators, the wavelengths $\lambda_{(0-15)}^{(1)}$ arrive at output port 1, $\lambda_{(16-31)}^{(1)}$ at output port 2, $\lambda_{(32-47)}^{(1)}$ at output port 3 and $\lambda_{(48-63)}^{(1)}$ at output port 0. This is repeated for input port 2 and input port 3 to create a 4 × 4 switching functionality device.

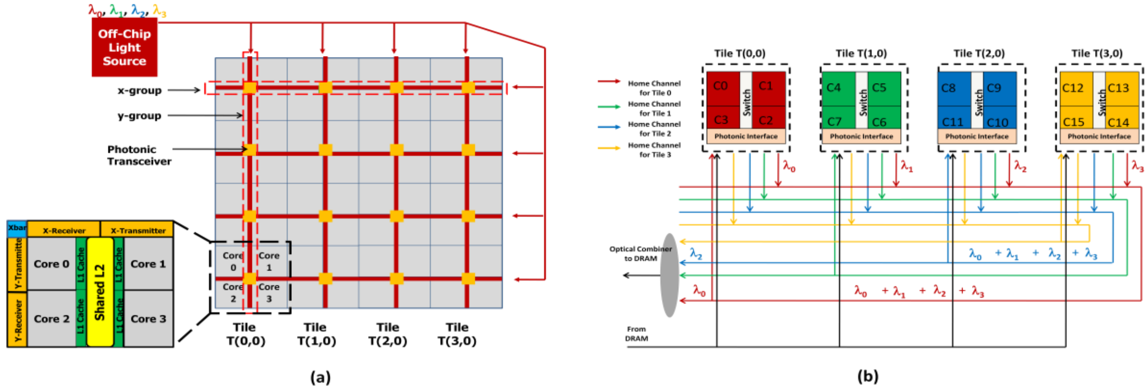

Figure 1: (a) Proposed layout of PROPEL architecture for 64 cores (cluster 0) and (b) the routing and wavelength assignment proposed for PROPEL for the x- dimension.

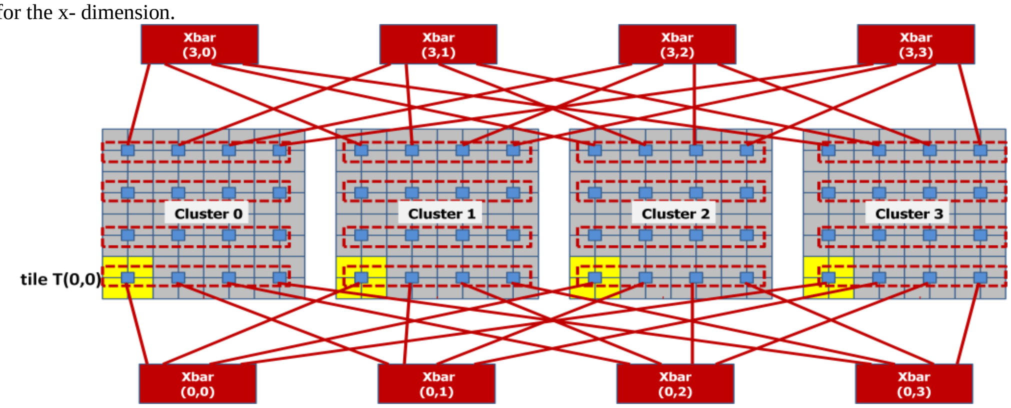

Figure 2: The proposed E-PROPEL architecture only top and bottom crossbars are shown for clarity.

Figure 2(b) shows the single micro-ring resonator implementation of a 64-wavelength 4×4 crossbar. It should be noted that this crossbar implementation is an extended version of the optical crossbar proposed by [4]. As shown in the figure, each waveguide is routed in a manner that allows it to come in close proximity to the other three waveguides. At these close proximity points a select range of wavelengths are switched between waveguides. This switching of light between two different waveguides allows light coming from one waveguide to be placed on another waveguide. In terms of functionality, this allows a single tile the ability to communicate with multiple tiles using only one input waveguide. Our results indicate, the signal ring crossbar consumes 1.38 mm^2 and has an optical loss of -4.32 dB.

For a further understanding of the single ring crossbar we will show how light from input 0 is switched and arrives on the 4 output waveguides. As light travels from input 0, it first encounters the intersection with the waveguide for input 1. At this point $\lambda^{(0)}_{(32-47)}$ is placed on the waveguide for input 1 allowing input 0 to communicate with output 2 using $\lambda^{(0)}_{(32-47)}$. Then light traveling from input 0 encounters the intersection with waveguide for input 3. At this point, $\lambda^{(0)}_{(0-15)}$ is placed on the waveguide for input 3 allowing input 0 to communicate with output 0 using $\lambda^{(0)}_{(0-15)}$. As the light continues traveling, it encounters the intersection with waveguide for input 2. At this point $\lambda^{(0)}_{(16-31)}$ is placed on the waveguide for input 2 allowing input 0 to communicate with output 1 using $\lambda^{(0)}_{(16-31)}$. Lastly $\lambda^{(0)}_{(48-63)}$ arrives at output 3 as it was the only light not switched which allows input 0 to communicate with output 3 using $\lambda^{(0)}_{(48-63)}$. This concept is expanded for other inputs to create a 64-wavelength 4 × 4 crossbar.

Figure 2(c) shows the double micro-ring resonator implementation of a 64-wavelength 4×4 crossbar. It should be mentioned that each micro-ring resonator in the figure represents 16 micro-ring resonators and is not shown for clarity. In the double rings crossbar, wavelengths are switched at locations where two waveguides are running parallel to each other. This results in an operating principle and functionality of the double rings crossbar that is identical to the single ring resonator crossbar. Our results indicate, the double rings crossbar consumes 0.0576 mm^2 and has an optical loss of -6.23 dB. In comparing the two designs, the single ring crossbar should be used in power sensitive designs as it is -1.92 dB less loss than the double rings design. The double rings crossbar should be used if area overhead is more important as it only consumes 0.057 mm^2.

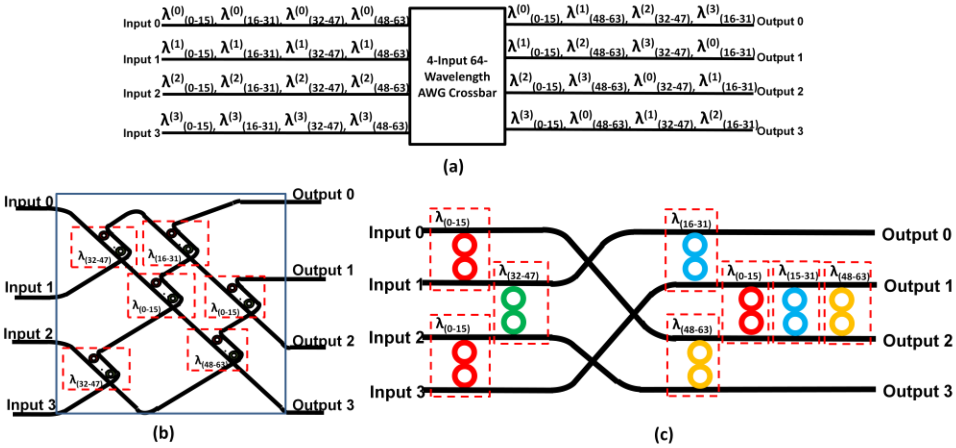

Figure 2:(a) functionality of each crossbar, (b)single ring crossbar and (c) double ring crossbar

5. Results

We simulated E-PROPEL on several traces including Uniform Random, and permutation patterns, such as Bit Reversal, Butterfly, Matrix Transpose, Complement and Perfect Shuffle using the cycle-accurate simulator OPTISIM in terms of saturation throughput and power [5]. Figure 4(a) shows the saturation throughput for the given traffic traces relative to the mesh network and Figure 4(b) shows the power dissipation relative for given traffic traces relative to the mesh network. Our results indicate that E-PROPEL occupies 3× less power and increases performance by 2× when compared to mesh topology [3].

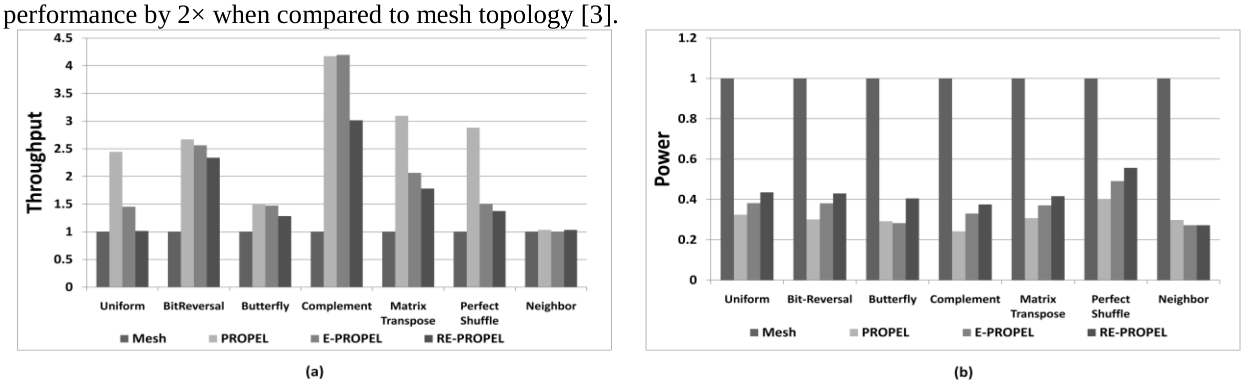

Figure 4: (a) Saturation throughput for select networks and (b) power dissipation for select networks.

6. Conclusion

In this paper, we propose two power and area efficient nanophotonic crossbars. Using the nanophotonic crossbars we implemented a 256-core NoCs called E-PROPEL. E-PROPEL combines 4 PROPEL networks by using the proposed nanophotonic crossbars. This creates an area and power efficient network topology which consumes 3× less power and increases performance by 2× when compared to mesh topology.

Acknowledgement: This research was sponsored by NSF grant CCR-0538945, ECCS-0725765 and CCF-0953398.

7. References

[1] G. Raymond and et. al., “Nanoelectronic and nanophotonic interconnect,” *Proceedings of the IEEE*, vol. 96, no. 2, pp. 230–247, Feb. 2008.

[2] A. Shacham, K. Bergman, and L. P. Carloni, “Photonic networks-on-chip for future generations of chip multiprocessors,” *in IEEE Transactions on Computers*, vol. 57, no. 9, pp. 1246–1260, September 2008.

[3] A. Kodi and R. Morris, “Design of a Scalable Nanophotonic Interconnect for Future Multicores,” *in ACM/IEEE Symposium on Architectures for Networking and Communications Systems (ANCS'09)*, Princeton, New Jersey, October 19-20, 2009.

[4] L. Zhou and et at., “Design and evaluation of an arbitration-free passive optical crossbar for on-chip interconnection networks,” *Applied Physics A*, vol. 95, no. 4, pp. 1111-1118, Oct 2008.

[5] A. Kodi and A. Louri, “Optisim: A System Simulation Methodology for Optically Interconnected HPC Systems,” *IEEE Micro*, vol. 28, no. 5, pp. 22-36.Sept/Oct 2008