

Dynamic Reconfiguration of 3D Photonic Networks-on-Chip for Maximizing Performance and Improving Fault Tolerance

Randy Morris†, Avinash Kodif and Ahmedouri ‡
 School of Electrical Engineering and Computer Science, Ohio University †
 Department of Electrical and Computer Engineering, University of Arizona ‡
 E-mail: kodi@ohio.edu, lourie@email.arizona.edu

ABSTRACT

As power dissipation in on-chip network is projected to be a major bottleneck, researchers are actively engaged in developing alternate power-efficient technology solutions. Nanophotonics along with 3D stacking are two disruptive technologies that are capable of delivering the communication bandwidth at low power dissipation. In this poster, we propose to combine nanophotonics and 3D stacking to develop a scalable, reconfigurable, and power-efficient interconnect called **R-3PO (Reconfigurable 3D Photonic Network-on-Chips)**. Our results indicate that the performance can be improved by 10%-25% for Splash-2, PARSEC and SPEC CPU2006 benchmarks along with saving of 6%-36% in terms of energy when compared to state-of-the-art on-chip electrical and optical networks.

INTRODUCTION

While Network-on-Chips (NoCs) design paradigm offers modular and scalable performance, increasing core counts leads to increase in serialization latency and power dissipation [1-3]. Emerging technologies such as nanophotonic interconnects (NIs) and 3D stacking are under serious consideration for meeting the communication challenges posed by the multicore due to their lower power dissipation and higher bandwidth over metallic interconnects [1-3].

In this poster, we leverage the advantages of two emerging technologies, NIs and 3D stacking, to design a high-bandwidth, low-latency, multi-layer, reconfigurable network, called **R-3PO (Reconfigurable 3D Photonic Network-on-chip network)**. R-3PO consists of 16 decomposed NI based crossbars placed on four optical communication layers, thereby eliminating waveguide crossing and reducing the optical power losses. In addition, we propose a reconfiguration algorithm whose purpose is to improve performance (throughput, latency) and bypass channel faults by adapting available network bandwidth. Our results indicate that the performance can be further improved by 10%-25% for Splash-2, PARSEC and SPEC CPU2006 benchmarks along with saving of 6%-36% in terms of energy.

R-3PO ARCHITECTURE

The R-3PO architecture consists of 256 cores in 64 tile configuration on a 400 mm² 3D IC. As shown in Figure 1, 256 cores are mapped on a 8×8 network with a concentration factor of four, called a tile. From Figure 1(a), the bottom layer, called the electrical die contains the cores, caches and memory controllers. The upper die, called the optical die, consists of the electro-optic transceivers layer and four decomposed nanophotonic crossbar layers. The electrooptic layer consists of all the front-end system drivers and the back-end receiver circuitry for photonics.

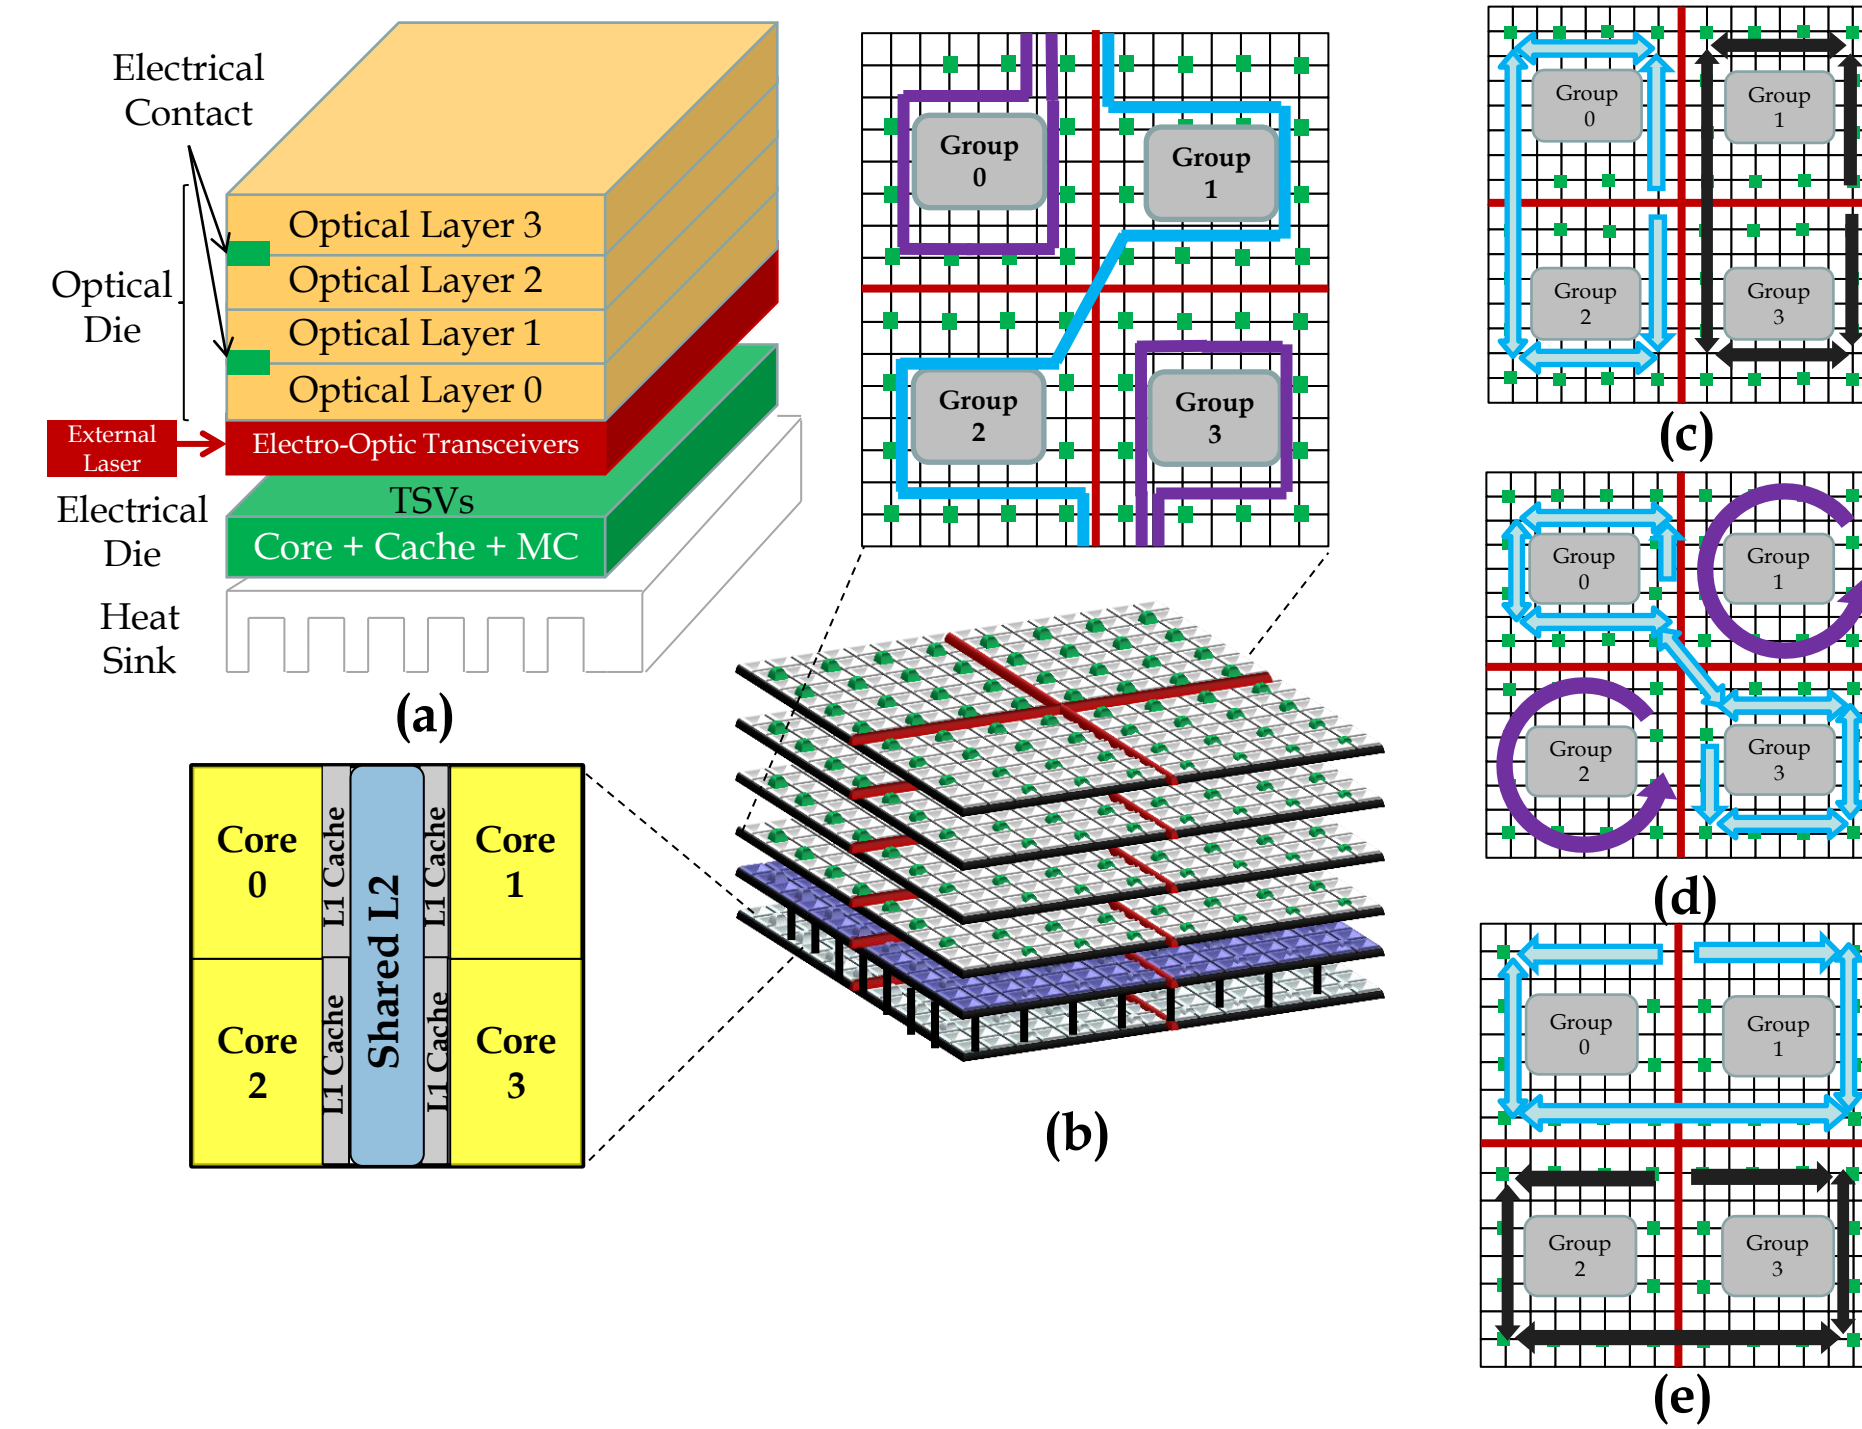


Figure 1: Proposed 256-core 3D chip layout. (a) Electrical die consists of the core, caches, and the memory controllers. The optical die on the lower most layer contain the electro-optic transceivers and four optical layers. (b) 3D chip with four decomposed nanophotonic crossbars with the top inset showing the communication in one layer 0 optical. (c) optical layer 1, (d) optical layer 2 and (e) optical layer 3.

RECONFIGURATION

The objective of reconfiguration is to improve performance by re-allocating bandwidth from under-utilized to over-utilized links. The design space of reconfiguration is large as there can be several combinations across multiple layers. Figure 2 shows four possible combinations that we evaluate as they cover most of the design space. Row-column matrix indicates the statically allocated communication. Figure 3 shows an example of reconfiguration between Group 0 and Group 3.

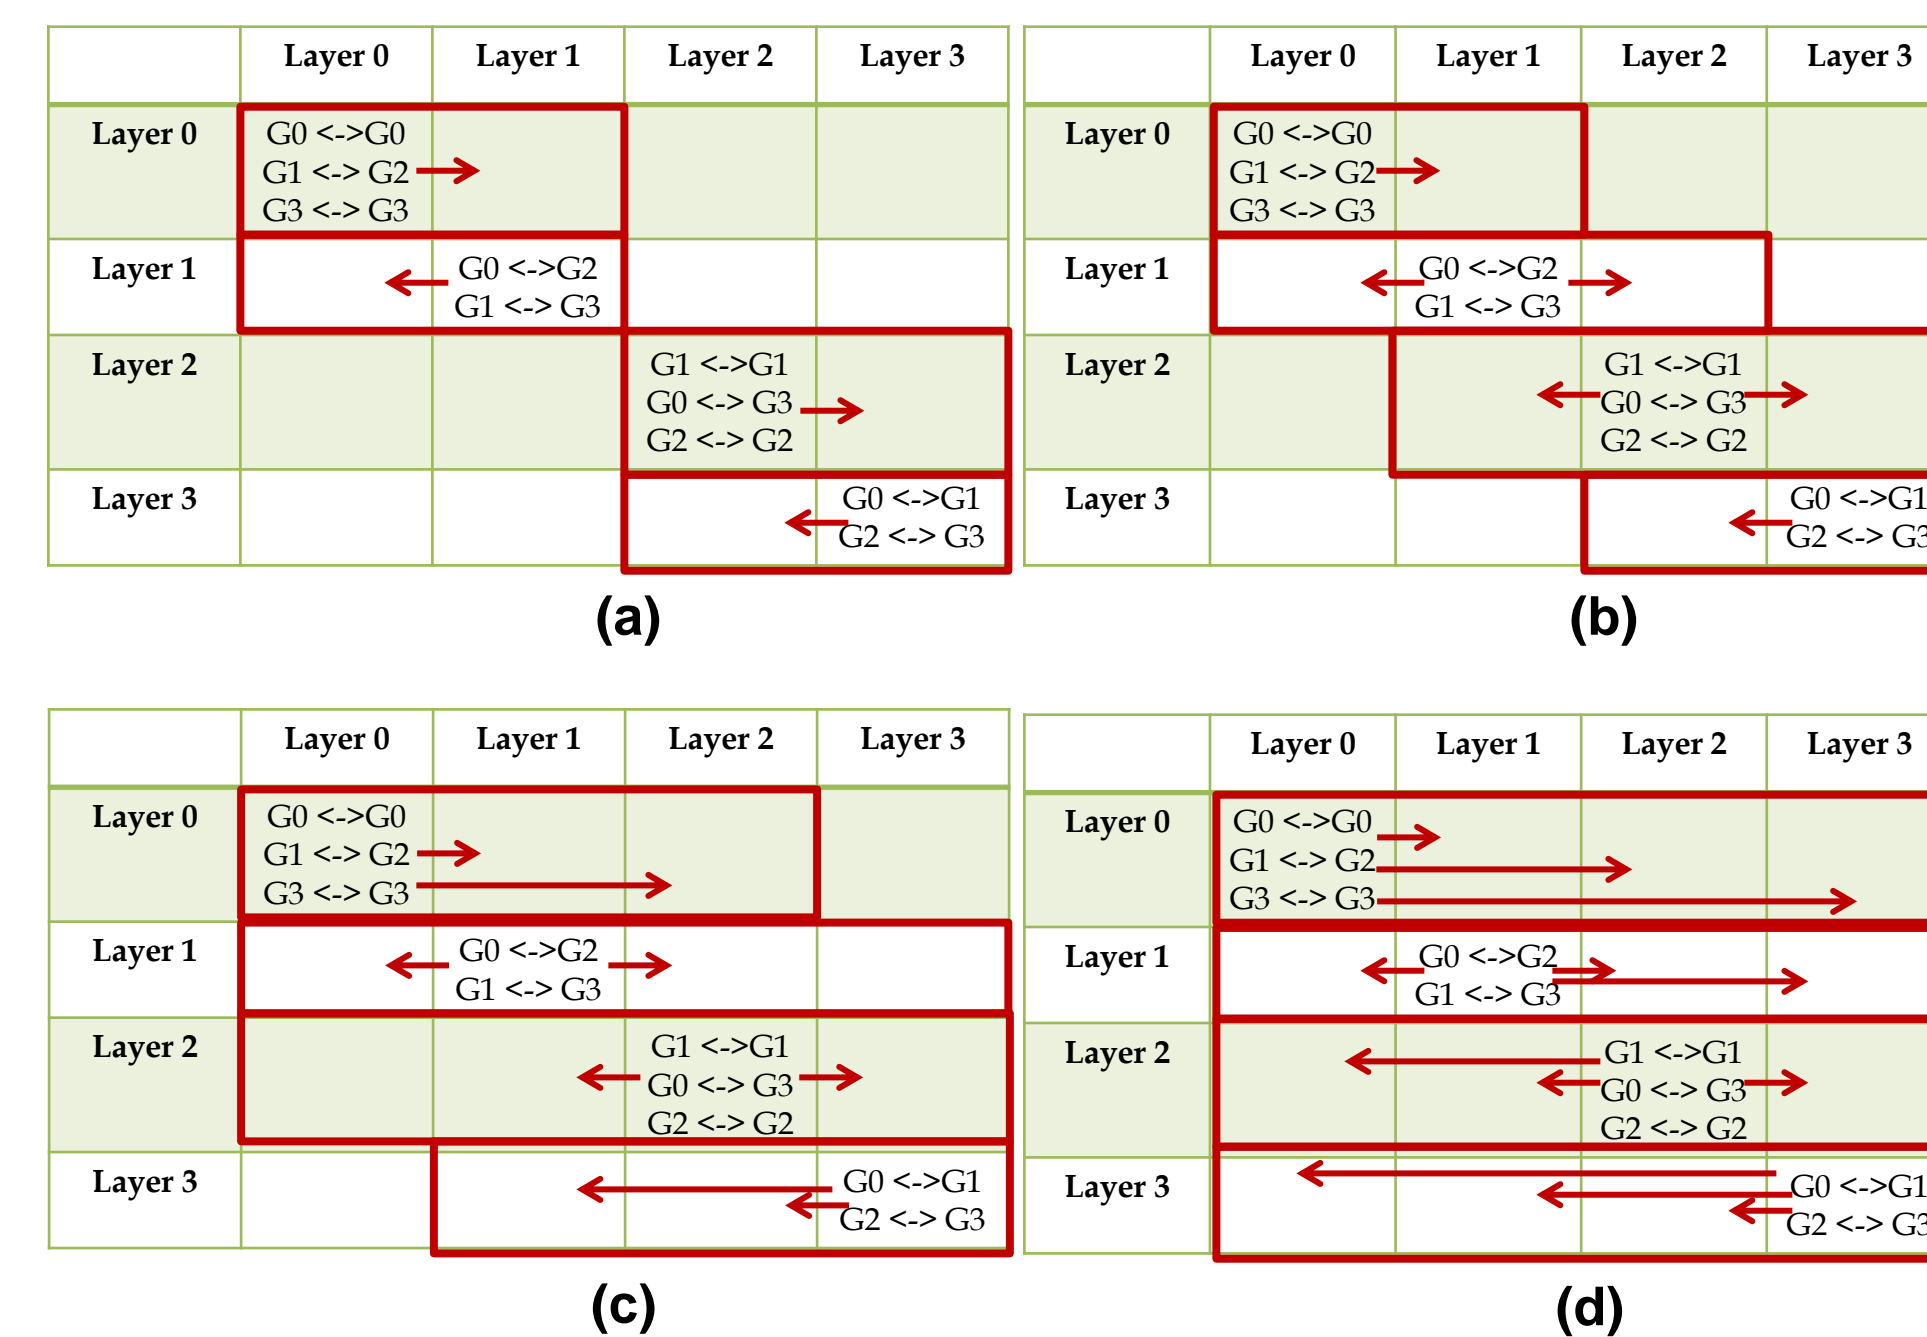


Figure 2: Various configurations evaluated: (a) R-3PO-L1, (b) R-3PO-LA, (c) R-3PO-L2 and (d) R-3PO-L3

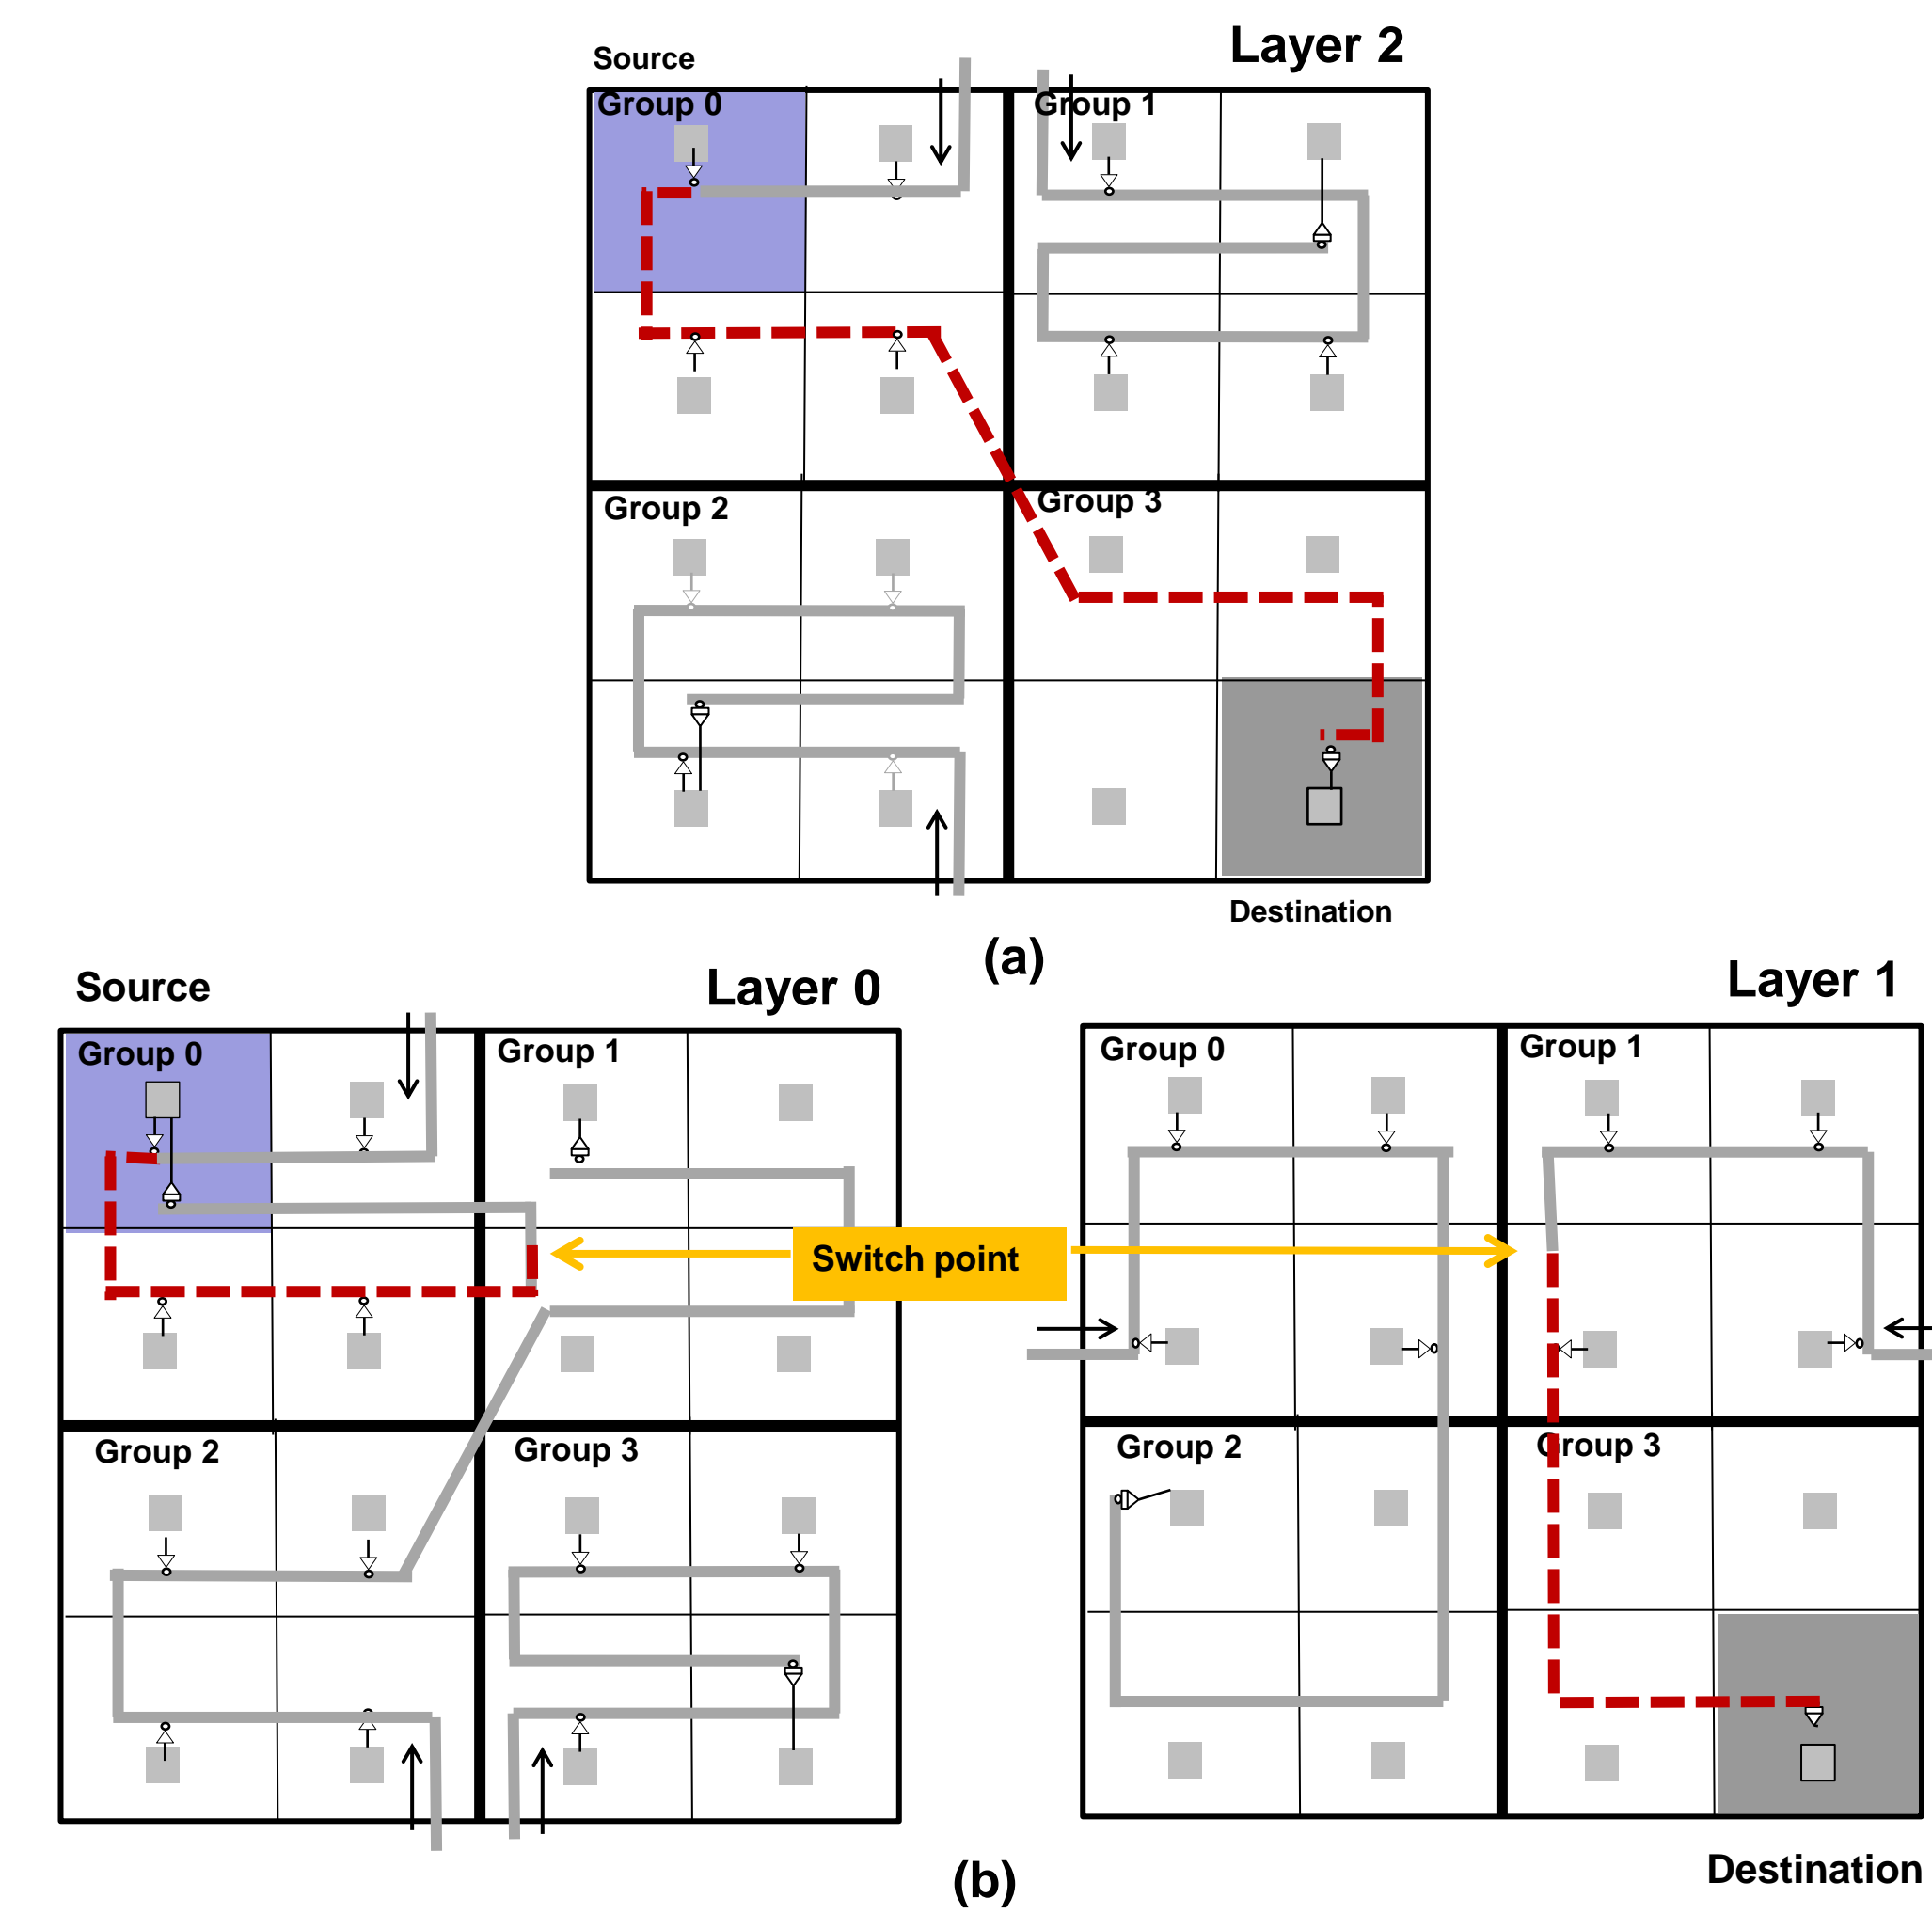


Figure 3: (a) Static communication between the source in Group 0 and destination in Group 3. (b) Illustration of reconfiguration between Groups 0 and 3 using partial waveguides from layers 0 and 1.

FAULT TOLERANCE

Fault tolerance allows data from a faulty channel to be switched to an adjacent layer that communicates with the same destination. Figure 4 shows an example of how fault tolerance is implemented.

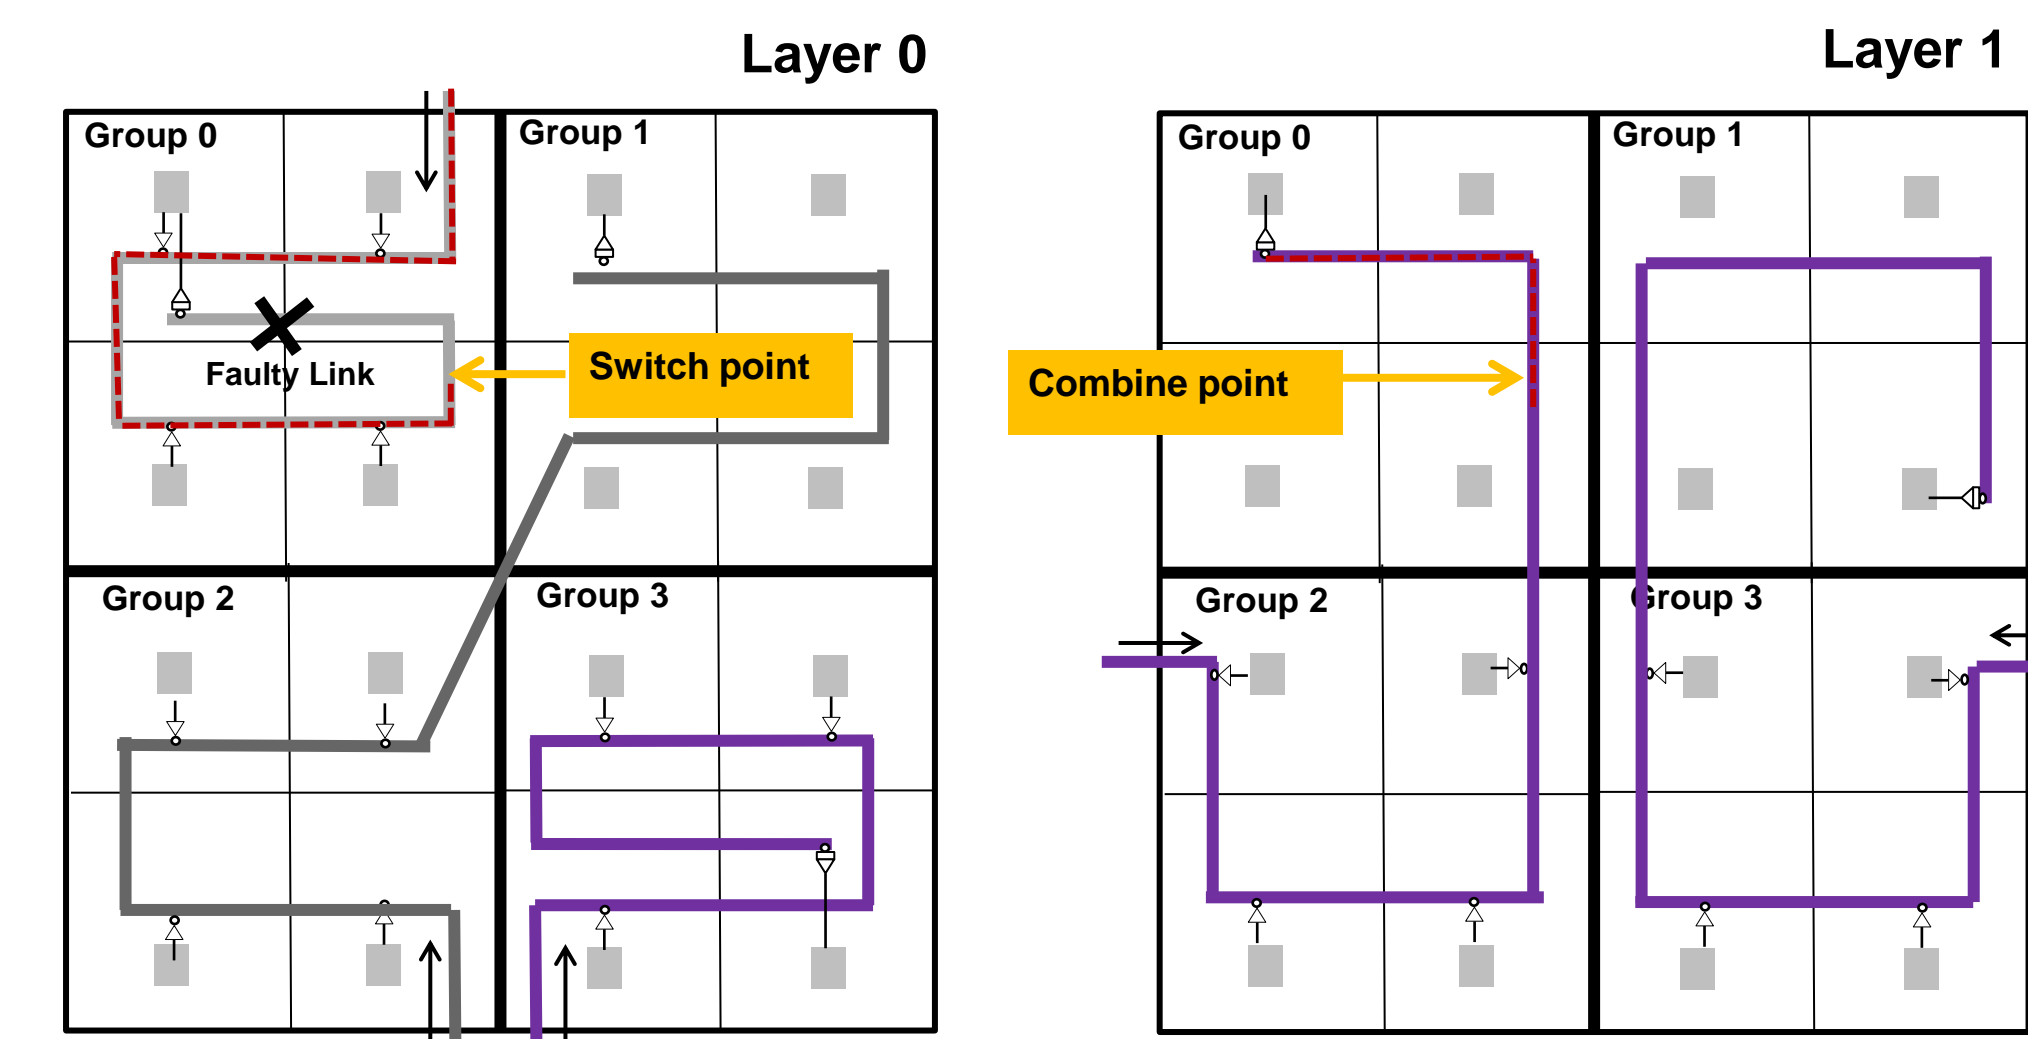


Figure 4: Fault tolerance in R-3PO.

RESULTS

Figure 5 shows the performance results for R-3PO compared to leading electrical (mesh & flattened butterfly) and optical (Corona, Firefly & MPNoC) networks. In Figure 5(a), R-3PO is one of the most power efficient topology with 36% less power than mesh. In Figure 5(b), R-3PO is the most performing network with a 2.5x increase over mesh. In Figure 5(c), R-3PO has the best energy-delay product.

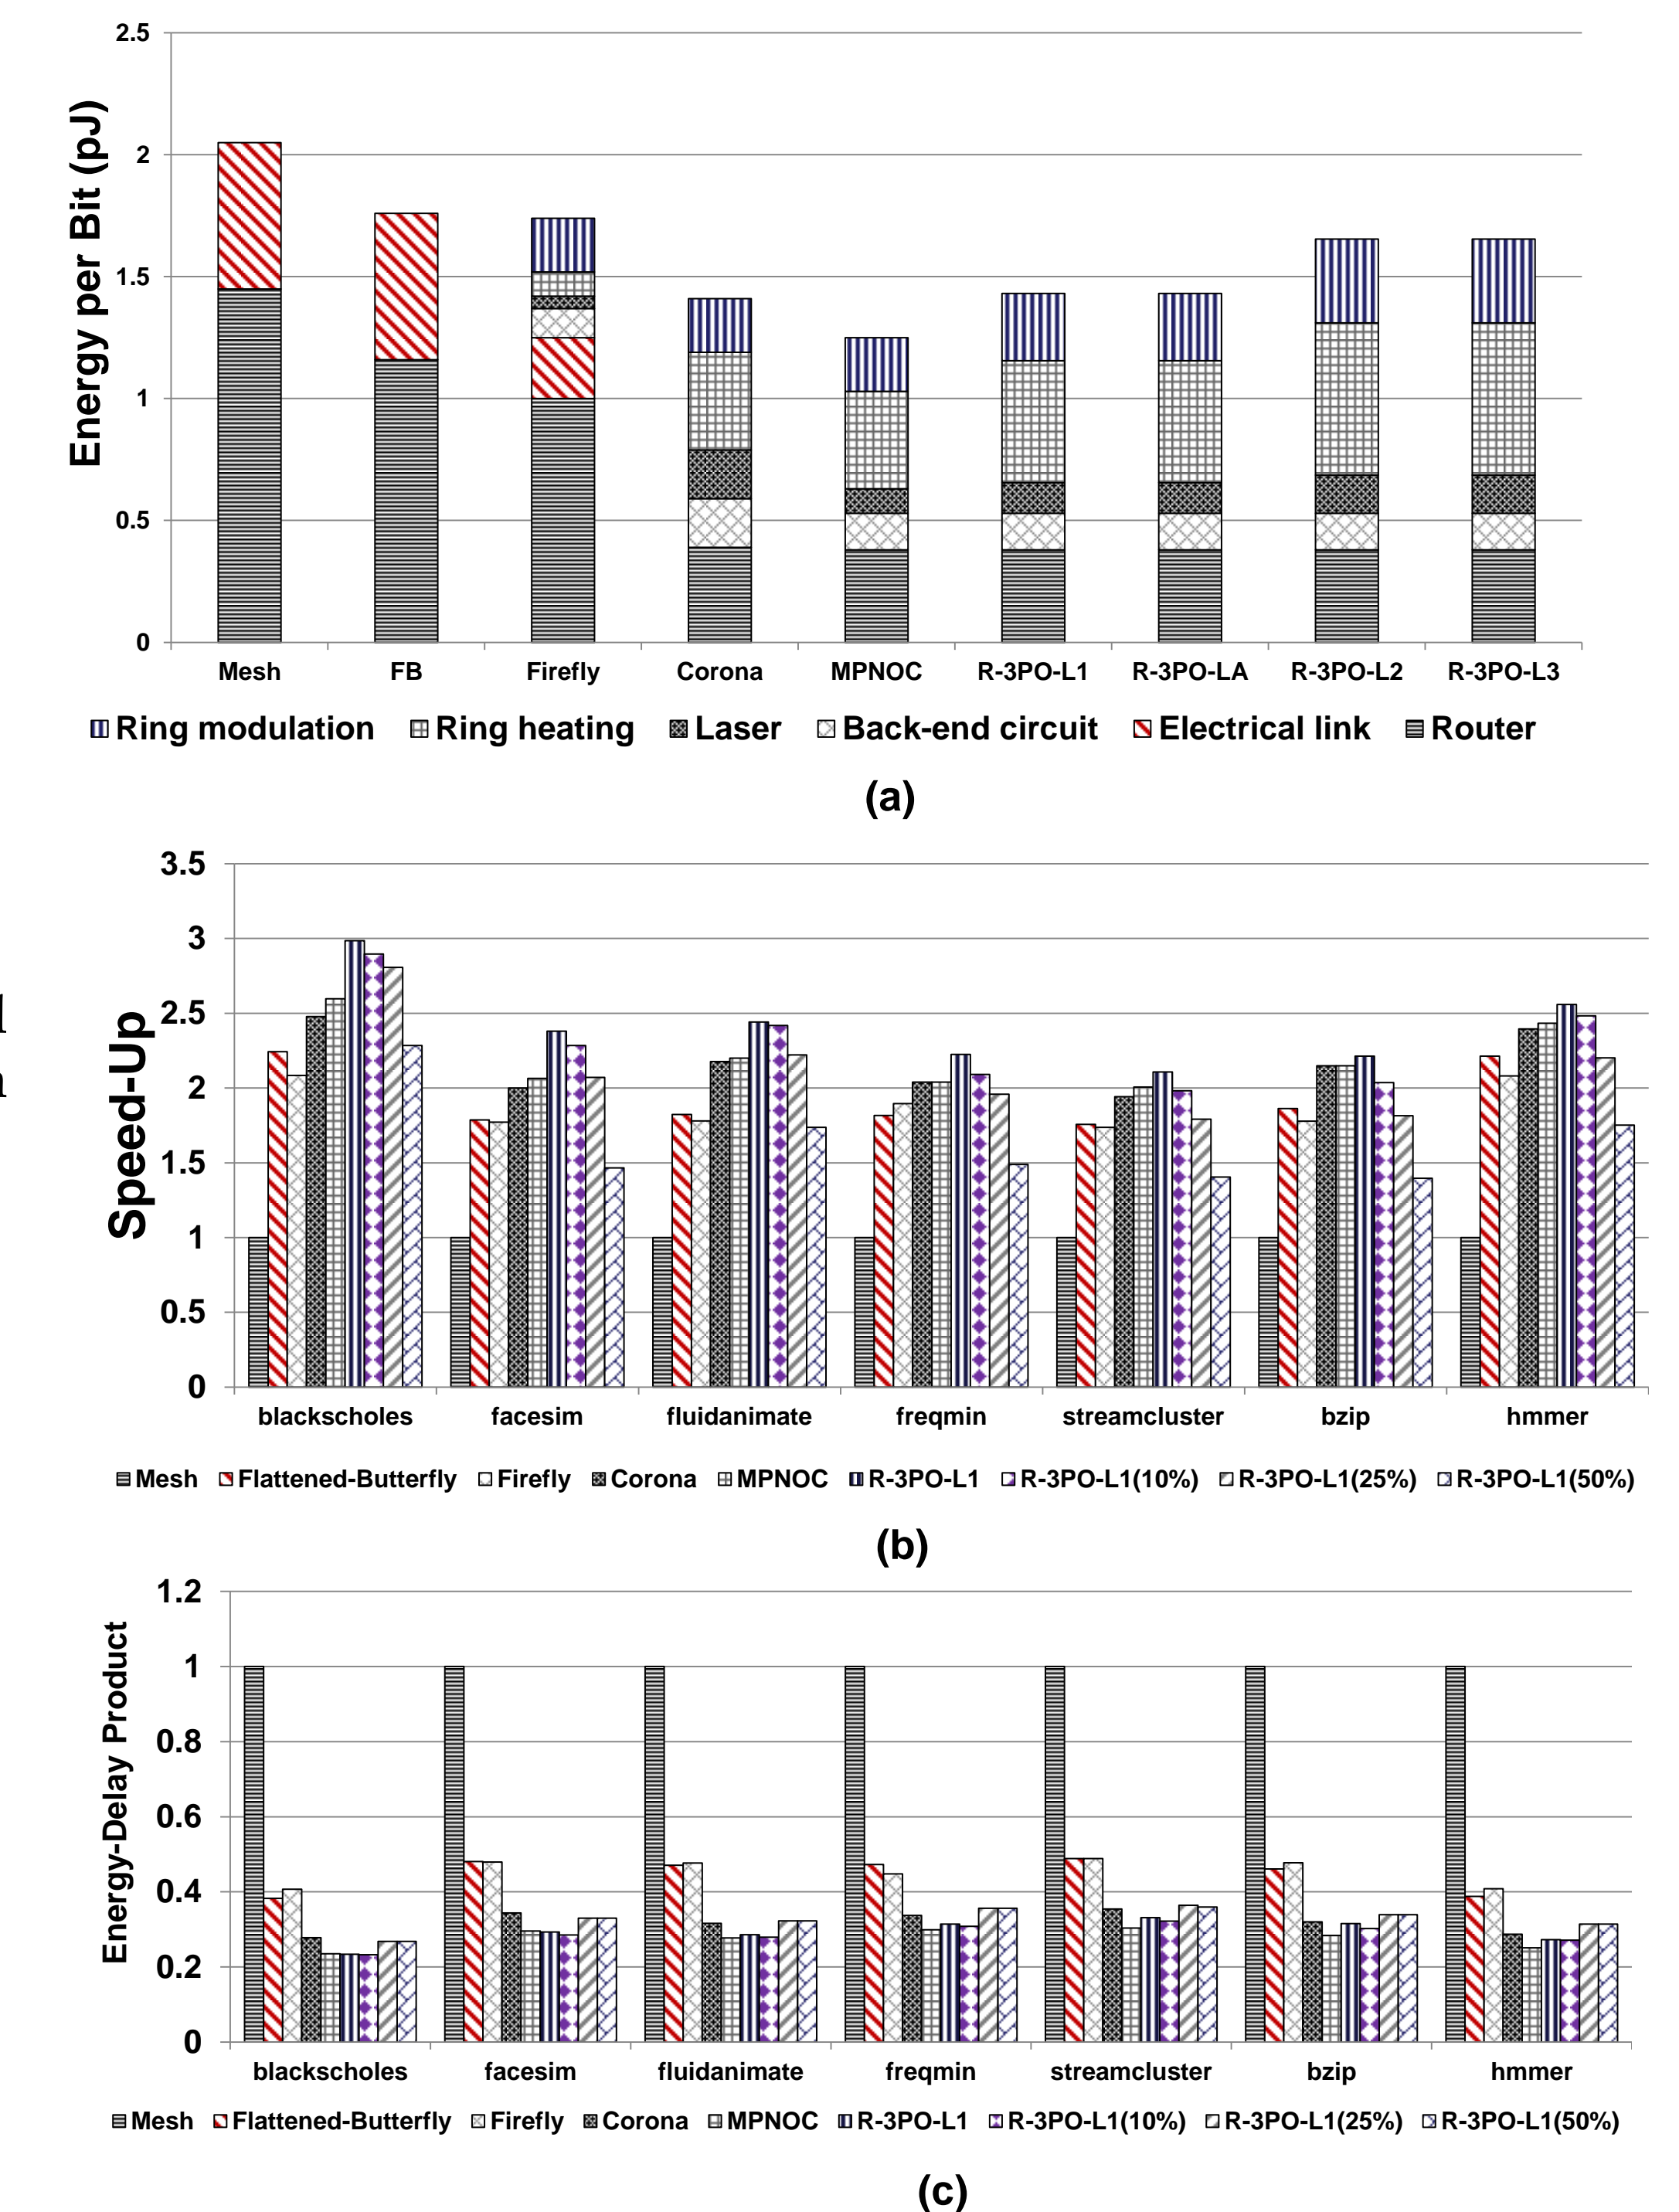


Figure 5: (a) Energy per bit, (b) performance results including degrade in performance from percentage of faults, and (c) energy-per bit results.

REFERENCES

- [1] Vantrease et. al., "Corona: System implications of emerging nanophotonic technology," in Proceedings of the 35th International Symposium on Computer Architecture, June 2008.
- [2] Y. Pan and et. al., "Firefly: Illuminating future network-on-chip with nanophotonics," in the Proceedings of the 36th Annual International Symposium on Computer Architecture, 2009.
- [3] X. Zhang and A. Lourie, "A multilayer nanophotonic interconnection network for on-chip many-core communications," in Proceedings of the Design and Automation Conference, 2010.