

Performance Evaluation for Three-Dimensional Networks-On-Chip

Brett Feero and Partha Pratim Pande

School of Electrical Engineering & Computer Science
Washington State University, PO BOX 642752, Pullman, Washington
(bfeero, pande)@eecs.wsu.edu

Abstract

Three dimensional (3D) integrated circuits (ICs) are capable of achieving better performance, functionality, and packaging density compared to more traditional planar ICs. On the other hand, Networks-on-Chip (NoCs) are an enabling solution for integrating large numbers of embedded cores in a single die. 3D NoC architectures combine the benefits of these two new domains to offer an unprecedented performance gain. In this paper, we develop a consistent and meaningful evaluation methodology to evaluate the performance of a variety of 3D NoC architectures compared to existing 2D counterparts. We demonstrate that the 3D NoCs are capable of achieving higher throughput, lower latency, and lower energy dissipation at the cost of small silicon area overhead.

1. Introduction and Motivation

With shrinking geometries, global interconnects are becoming the principal performance bottleneck for high-performance Systems-on-Chip (SoCs) [1] [2]. These long interconnects are quickly becoming a performance impediment in terms of communication latency and power. The Network-on-Chip (NoC) model is emerging as a revolutionary methodology in solving the performance limitations arising out of long interconnects. In addition to providing a solution for the global wire delay problem, the NoC paradigm also eases integration of high numbers of intellectual property (IP) cores in a single SoC. However, the conventional two dimensional (2D) IC has limited floor-planning choices, and consequently it limits the performance enhancements arising out of NoC architectures. As shown in Figure 1, three-dimensional (3D) integrated circuits (ICs), which contain multiple layers of active devices, have the potential for enhancing system performance [3]. According to [3], three-dimensional ICs allow for performance enhancements even in the absence of scaling. This is the result of each transistor being able to reduce interconnect length and access more nearest neighbors. Besides this clear




Figure 1. Three-dimensional integrated circuit

benefit, package density is increased significantly, power is reduced from shorter wires, and circuitry is more immune to noise [3]. The performance improvement arising from the architectural advantages of NoCs will be significantly enhanced if 3D ICs are adopted as the basic fabrication methodologies. The amalgamation of two emerging paradigms, NoC and 3D IC allows for the creation of new structures that enable significant performance enhancements over more traditional solutions.

In [4], 3D ICs were proposed to improve performance of chip multi-processors. Drawing upon 3D IC research, they chose a hybridization of busses and networks to provide the interconnect fabric between CPUs and L2 caches. This fusion of NoC and bus architectures was evaluated for performance using standard CPU benchmarks. However, this analysis pertains only to chip multi-processors and does not consider the use of three-dimensional network structures for application-specific SoCs. Some of these concerns are addressed in [5], which compares 2D MESH structures with their 3D counterparts by analyzing the zero-load latency of each architecture. This is an evaluation that shows some of the advantages of 3D NoCs, but it neither applies any real traffic pattern, nor does it measure other relevant performance metrics. We endeavor to address these concerns by applying real traffic patterns in a cycle-accurate simulation, and measuring performance through established metrics [6] for 3D NoC structures.




Figure 2. NoC architectures: (a) 2D MESH, (b) 2D Folded Torus, (c) 3D MESH, (d) 3D Stacked MESH

In this paper, we will establish that three-dimensional NoCs are a viable and advantageous alternative to current interconnect structures. We will introduce multiple 3D network topologies and compare their performance with respect to some proven 2D topologies using realistic traffic patterns in terms of standard metrics [6].

2. 3D NoC Architectures

Designing in the third dimension provides new challenges in the form of a new degree of freedom. In this paper, we introduce four different 3D NoC architectures. The performance of these 3D NoCs will be evaluated with respect to more traditional 2D structures. The 2D architectures considered are the MESH and Folded Torus [6].

2.1 3D MESH and Torus

The first architecture considered here is the 3D MESH. The three-dimensional MESH (shown in Figure 2c) is simply an extension of the popular planar structure (Figure 2a). In case of a system with 64 IP blocks, the 3D NoC consists of four blocks in each dimension (4x4x4). It will employ 7-port switches as shown in Figure 3a (1 each for North, South, East, West, up, down, and the IP). In previous works [4], the 3D MESH structure is discarded due to contention issues, exacerbated power consumption, and multiple hops while traversing the z -dimension. However,




Figure 3. Switches: (a) 3D MESH, (b) 3D Stacked MESH

we will show in this paper that this structure is clearly a viable solution as an interconnect fabric.

A similar structure to MESH is a torus. Toroidal structures differ in that the switches at the edges wrap around to the opposite side. This ensures that all switches have equal numbers of ports. However, this leads to long wrap-around wires, so the Folded Torus architecture (Figure 2b), in which all wires are the same length, was designed as a mitigating solution [6]. From this architecture, a three-dimensional Folded Torus similar to Figure 2c is easily extensible. This is another architecture considered in our analysis.

2.2 3D Stacked MESH

Stacked MESH is a NoC structure first shown in [4]. It features multiple layers of 2D MESH structures stacked one upon another, with communication in the third dimension (as

shown in Figure 2d) taking place through busses instead of a packet-switched network. This is desirable due to the short distance between the 2D MESH layers. Similar to Stacked MESH, Stacked Torus is a slight modification where folded tori are stacked instead of MESH structures. This design is considered promising due to the nature of the 3D silicon structure. Layers of silicon are stacked upon each other, leading to interlayer distances that are much less than the lateral inter-switch lengths. Vertical links may be around 10 μm and certainly less than 100 μm [3]. On the other hand, lateral links can easily be more than 1mm [7]. Furthermore, the z -dimension will not scale proportionally to the x - and y -dimensions. So, due to the reduced capacitances (shorter wires and fewer links), busses are reasonable solutions for interlayer communication.

A typical switch for stacked architectures will contain one link to the IP block, 2 links in the x -dimension, 2 links in the y -dimension, and one link to the bus, as shown in Figure 3b. This eliminates one of the ports for the z -direction compared to a standard 3D MESH or torus, reducing overall power consumption and simplifying routing. However, at a certain size in the vertical dimension, this structure will cease to see speed benefits due to contention issues over the bus. In [4], it is stated that contention issues should not be prohibitive until 9 layers of silicon are used. In order to achieve maximum performance from this 3D NoC/bus hybrid, and appropriate bus must be chosen, so to be consistent with [4], a dynamic Time-Division Multiple Access (dTDMA) bus is used here.

3. Performance Metrics

In order to properly analyze different network-on-chip architectures, a standard set of metrics must be used [6]. We will compare different topologies on four benchmarks: throughput, latency, energy and area overhead. We consider the same performance metrics as shown in [6].

3.1 Throughput

Throughput is a measure of how much data can be transmitted across a network. For purposes of a message-passing system, throughput T is given by the equation

$$T = \frac{(\text{Total Messages Completed}) \times (\text{Message Length})}{(\text{Number of IP Blocks}) \times (\text{Time})} \quad (1)$$

Total messages completed are the number of messages which successfully traverse the network from source to destination. *Message length* refers to the number of flits a message consists of, and *Number of IP Blocks* signifies the number of Intellectual Property blocks that send data over the network. *Time* is length of time in clock cycles between the generation of the first packet and the reception of the last. It can be seen that throughput is measured in flits/IP block/cycle, where a throughput of 1 signifies that every IP block is accepting a flit in each clock cycle. Accordingly,

throughput is a measure of the maximum amount of sustainable traffic. Throughput will be dependent on a number of parameters including the number of links in the architecture, the average hop count, the number of ports per switch, and injection load. Injection load is measured by the number of flits injected in to the network per IP block per cycle. Consequently it has the same unit as the throughput.

3.2 Latency

Latency is defined as the time in clock cycles elapsed from the transfer of the header flit by the source IP to the acceptance of the tail flit by the destination IP block. Latency is characterized by three delays: sender overhead, transport latency, and receiver overhead.

$$L_i = L_{\text{sender}} + L_{\text{transport}} + L_{\text{receiver}} \quad (2)$$

Flits must traverse a network while traveling from source to destination. With different routing algorithms and switch architectures, each packet will experience a unique latency. As a result, network topologies will be compared by average latency. Let P be the number of packets received in a given time period, and let L_i be the latency of the i th packet. Average latency is therefore given by the equation:

$$L_{\text{avg}} = \frac{\sum_{i=1}^P L_i}{P} \quad (3)$$

3.3 Energy

When a flit travels across an interconnection network, both the logic gates of the switches and the interconnect wires toggle. As a result, energy is dissipated in each step. Flits travel across a network in hops from switch to switch, and as such, we model energy at the most basic level as energy per flit per hop.

$$E_{\text{hop}} = E_{\text{switch}} + E_{\text{wire}}, \quad (4)$$

where E_{switch} and E_{wire} are the energy dissipated by each switch and inter-switch wire segments respectively.

The energy of a packet of length n flits that completes h hops is given by

$$E_{\text{packet}} = n \sum_{j=1}^h E_{\text{hop},j} \quad (5)$$

In order to characterize different network topologies, we are interested in comparing the average energy dissipated per packet while communicating between two IP blocks. If P packets are transmitted then the average energy dissipated per packet is given as

$$\overline{E_{\text{packet}}} = \frac{\sum_{i=1}^P E_{\text{packet},i}}{P} = \frac{\sum_{i=1}^P \left(n_i \sum_{j=1}^{h_i} E_{\text{hop},j} \right)}{P} \quad (6)$$

3.4 Area

Finally, to appropriately characterize different networks, we must consider the silicon area overhead of each

Table 1. Simulation parameters

Topology	IPs	Buffer Depth (flits)	Port Count	Message Length (Flits)
2D MESH/Torus	64	2	5	64
Stacked MESH/Torus	64	2	6	64
3D MESH/Torus	64	2	7	64

implementation. As the switches form an integral part of the infrastructure, it is important to determine the amount of relative silicon area they consume. Additionally area overhead arising out of inter-switch wire segments and routing of those wires need to be considered [6]. As a first step in this work, we investigate only the switch area overhead introduced by 3D NoCs.

3.5 Simulation Methodology

To model different NoC structures, a cycle-accurate network simulator is employed. It is flit-driven and uses wormhole routing. In wormhole routing, the packet is divided into fixed length flow control units or flits. The header flit holds the routing and control information. It establishes a path, and subsequent payload or body flits follow that path.

The simulator models traffic by a self-similar distribution. In the past, Poisson distributions have been used to simulate network traffic [6]; however, a self-similar distribution has been shown to be a more accurate and realistic model [8]. Self-similar traffic is characterized by modeling injections through a multitude of ON-OFF message sources according to a Pareto distribution. The network switches, dTDMA arbiter, and FIFO buffers were modeled in VHDL. Although the simulator is capable of running with an arbitrary specification for each parameter, each switch was designed with 4 virtual channels per port and 2-flit deep virtual channel buffers as discussed in [6]. Synopsys Design Analyzer was used to synthesize the hardware description using a 90nm standard cell library from CMP [9].

The simulation is initially run for 1,500 cycles to allow the network to stabilize, and it is subsequently run for 18,500 more cycles. The simulator provides statistics for energy, throughput, and latency. In order to calculate energy, the simulator traces each hop and the associated inter-switch wire and switch energy. It tracks energies in each hop since not all links are congruent.

4. Experimental Results and Analysis

In accordance with the prescribed methods, each NoC architecture mentioned above were analyzed for the four performance metrics: throughput, latency, energy and area. To do this, we considered 64-IP versions of each of the

network topologies mapped into 400mm². In the case of the 2D architectures, we assume a 20mm × 20mm die, while we propose four 10mm × 10mm layers for the three-dimensional arrangements. It can clearly be seen that both the total silicon area and the interconnect lengths (in the *x*- and *y*-dimensions) are identical to those of a 20mm × 20mm two-dimensional model. These architectures were then simulated using the cycle-accurate simulator described in section 3.5. The simulation parameters are shown in Table 1.

4.1 Throughput

We now characterize the new architectures based on the throughput metric. Throughput is given by the number of accepted flits per IP per cycle. This metric, therefore, is closely related to the maximum amount of sustainable traffic in a certain network type. Figure 4a shows the variation of throughput with injection load for the NoC architectures under consideration here. A network cannot accept more traffic than is supplied, and limitations in routing and collisions cause saturation before throughput reaches unity. From Figure 4a, we see clearly that the three-dimensional solutions to a 64-IP structure are superior as far as throughput is concerned. This increase in throughput is related to three factors. The number of physical links, the average hops and the number of ports per switch. The number of links in a 2D MESH is given by the following equation

$$links_{2D_MESH} = N_1(N_2 - 1) + N_2(N_1 - 1), \quad (7)$$

where N_i represents the number of switches in the *i*-dimension. With an 8×8 2-D MESH-based NoC this yields 112 links. In a three-dimensional MESH the number of links is given as follows.


$$links_{3D_MESH} = N_1N_2(N_3 - 1) + N_1N_3(N_2 - 1) + N_2N_3(N_1 - 1) \quad (8)$$

Following (8) the number of links in a 4×4×4 3D MESH, turns out to be 144.

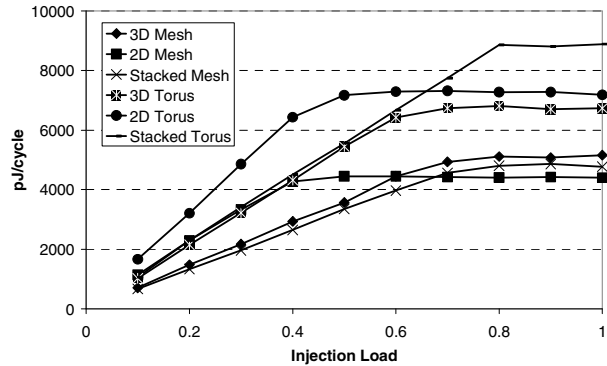
We see that there is a 29% increase in links, allowing for 29% more flits to exist in the network structure. However, this does not account for more than 70% increase in throughput. To characterize this, we must consider the average number of hops in the network. Following [5], the average number of hops in a MESH-based NoC is given by,

$$hops_{NoC} = \frac{n_1n_2n_3(n_1 + n_2 + n_3) - n_3(n_1 + n_2) - n_1n_2}{3(n_1n_2n_3 - 1)} \quad (9)$$


For our 4×4×4 3D MESH and 8×8 2D MESH, average hop counts are 3.81 and 5.33, respectively. There are 40%



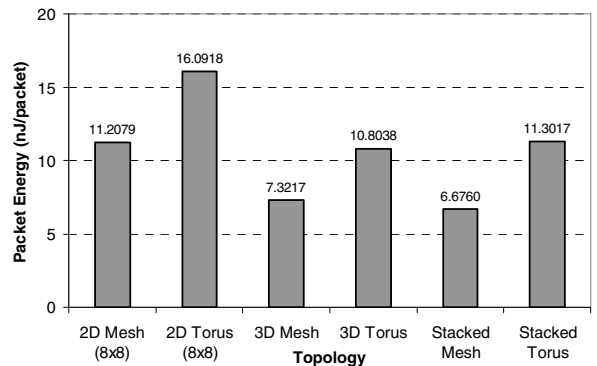
(a)



(c)



(b)



(d)

Figure 4. Experimental results: (a) Throughput vs. injection load, (b) Latency vs. injection load, (c) Cycle energy vs. injection load, (d) Packet energy

more hops in the 2D MESH compared to that in 3D MESH. Consequently flits in 3D MESH needs to traverse less number of stages between a pair of source and destination than the 2D counterpart. As a result of this we expect a corresponding increase in throughput. The average hop count for the stacked MESH should be slightly lower than the 3D MESH due to the bus spanning all layers of the chip, and the hop counts for the toroidal structures should be approximately equal to their MESH-based counterparts. Finally, the switch architecture needs to be taken into account. With the addition of two ports in the 3D MESH and one port in the Stacked MESH, number of flits transported through the switch increases. These three factors together account for the significant increase in throughput for the 3D structures.

4.2 Latency

Latency is a metric that is concerned with how quickly a message can be transferred between a pair of communicating blocks across a network. Figure 4b shows latencies for the architectures under consideration. Once again, we see that the 3D architectures outperform the 2D implementations.

This is a direct effect of fewer hops between nodes, i.e. a smaller zero-load-latency, and fewer contention issues.

4.3 Energy

In the System-on-Chip realm, energy dissipation characteristics of the interconnect structures are crucial, as the interconnect fabric can consume a significant portion of the overall energy budget. The energy dissipation profiles of the NoC architectures under consideration are shown in Figures 4c and 4d. Energy dissipation is largely dependent on two factors: architecture and injection load. We consider these two parameters as the independent factors in our analysis. As shown in equation (4), the energy dissipation in a NoC depends on the energy dissipated by the switch blocks and the inter-switch wire segments. Both these factors depend on the architectures. The design of the switch varies with the architecture and inter-switch wire length is also architecture dependent [7]. In addition to the network architecture, injection load has a clear effect on the total energy dissipation of a NoC. Intuitively, it is clear that with more packets traversing the network, power will increase. This is why packet energy is an important attribute for characterizing

NoC structures. It is clear from Figure 4d that it takes less energy for a 3D MESH to transmit a packet than a 2D MESH. However, Figure 4c reveals that, at saturation, a 2D MESH will dissipate less power than its three-dimensional counterpart. This is the result of the lower 2D MESH throughput, and the 3D MESH only consumes more energy because it can transmit more flits.

According to Figure 4d, the three-dimensional structures show significantly lower energy per packet than two-dimensional implementations. These results follow equation (9), with the smaller hop counts.

The introduction of a third dimension reduces the average hop count between two communicating IP blocks, as shown in [5]. Average packet energy is dependent on the number of hops from source to destination. Consequently, as the number of hops is reduced in a three-dimensional system, average packet energy drops accordingly.

4.4 Area

The final performance metric used in this study is the silicon area consumed by the switch blocks. The critical idea here is that area consumed by infrastructure cannot be utilized for functional units. The principal component of extra silicon area overhead arises from the switch blocks. We designed switch blocks corresponding to all the NoC architectures investigated in this work and synthesized them using 90 nm standard cell libraries from CMP [9]. Vertical vias in 3D NoCs use a non-trivial amount of floor area, but with via pitch near $2\mu\text{m}$ [3], the area used by vertical vias is nearly insignificant. The switch area overhead, represented in terms of equivalent 2-input NAND gates for each architecture is shown in Table 2.

The 3D structures utilize more silicon area as the switches require more ports. As shown in Figure 5, this is an easily-justified concession considering the lower energy dissipation and higher throughput provided by the 3D NoCs compared to their 2D counterparts.

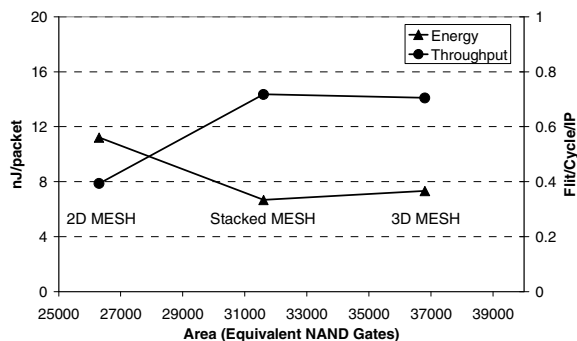


Figure 5. Throughput and energy vs. silicon area

Table 2. Switch Area Overhead

	2-input NAND gates
2D MESH/Torus	26,300
Stacked MESH/Torus	31,600
3D MESH/Torus	36,800

5. Conclusions

Three-dimensional networks-on-chip are natural extensions of two-dimensional designs. In this paper we have demonstrated that besides reducing the footprint in a fabricated design, three-dimensional network structures provide a better performance compared to traditional, 2D NoC architectures. Cycle-accurate simulations have validated this claim, showing that despite an area penalty, significant gains are made in terms of energy, throughput, and latency. The Network on chip (NoC) paradigm continues to attract significant research attention in both academia and industry. With the advent of 3D ICs the achievable performance benefits from NoC methodology will be more pronounced as shown in this paper. Consequently this will facilitate adoption of NoC as a mainstream design solution for big multi-core system chips.

6. References

- [1] L. Benini and G. De Micheli, "Networks on Chips: A New SoC Paradigm," *IEEE Computer*, Jan. 2002, pp. 70-78.
- [2] P. Magarshack and P.G. Paulin, "System-on-Chip beyond the Nanometer Wall," Proceedings of 40th Design Automation Conf. (DAC 03), ACM Press, 2003, pp. 419-424.
- [3] A. W. Topol et al., "Three-dimensional integrated circuits," *IBM J. RES. & DEV. VOL. 50 NO. 4/5 JULY/SEPTEMBER 2006*
- [4] F. Li et al., "Design and Management of 3D Chip Multiprocessors Using Network-in-Memory", Proceedings of the 33rd International Symposium on Computer Architecture (ISCA'06), pp. 130-141
- [5] V. F. Pavlidis and E. G. Friedman, "3-D Topologies for Networks-on-Chip", Proceedings of IEEE SOC Conference, (SOCC 2006), pp. 285-288.
- [6] P. P. Pande, C. Grecu, M. Jones, A. Ivanov, R. Saleh, "Performance Evaluation and Design Trade-offs for Network on Chip Interconnect Architectures", *IEEE Transactions on Computers*, Vol. 54, no. 8, August 2005, pp. 1025-1040.
- [7] C. Grecu, P. P. Pande, A. Ivanov, R. Saleh, "Timing Analysis of Network on Chip Architectures for MP-SoC Platforms", *Microelectronics Journal*, Elsevier, Vol. 36, issue 9, pp. 833-845.
- [8] K. Park and W. Willinger, *Self-Similar Network Traffic and Performance Evaluation*. John Wiley & Sons, 2000.
- [9] <http://cmp.imag.fr/index.php>