

2L-2D Routing for Buffered Mesh Network-on-Chip

Rose George Kunthara¹[0000-0001-5376-5718], Neethu K¹[0000-0001-5897-7146],
Rekha K James¹, Simi Zerine Sreeba²[0000-0002-2716-270X], Tripti S Warriar³,
and John Jose⁴[0000-0002-2558-0991]

¹ Division of Electronics, School of Engineering, CUSAT, Cochin, India
rosekunthara87@gmail.com, neethukuriyedam@gmail.com,
rekhajames@cusat.ac.in

² Viswajyothi College of Engineering and Technology, Muvattupuzha, India
simi.abie@gmail.com

³ Department of Electronics, CUSAT, Cochin, India
tripti@cusat.ac.in

⁴ Indian Institute of Technology Guwahati, India
johnjose@iitg.ac.in

Abstract. The rise in complexity and number of processing cores in SoC has paved way to the development of efficient and structured on-chip communication framework known as Network on Chip (NoC). NoC is embraced as an interconnect solution for the design of large tiled chip multiprocessors (TCMP). It is characterized by performance metrics such as average latency, throughput and power dissipation which depend on underlying network architecture. In this paper, we propose 2L-2D (Two Layer Two dimensional) architecture to enhance performance of conventional buffered 2D mesh NoC where two identical layers of 8x8 meshes are stacked one on top of the other. 2L-2D uses conventional 5-port virtual channel router (VCR) architecture and vertical interconnections are made by utilizing unused ports at edge routers only. Experimental results indicate that our proposed approach improves throughput and network saturation point whereas average flit latency and power dissipation is considerably reduced when compared with standard 5-port 2D mesh and torus designs.

Keywords: Network-on-Chip · Virtual Channel Router · Average latency · Throughput.

1 Introduction

Rapid progress and innovations in IC technology have led to massive rise in transistor integration which resulted in the evolution of complex System-on-Chip (SoC) comprising of IP (Intellectual Property) cores that are connected either by classical shared bus or point to point intercommunication architectures. Network-on-Chip (NoC), a packet-switched network, has emerged to overcome integration restrictions of SoC and interconnect associated issues like global

wire delay problem which arise due to technology scaling. NoC communication is widely employed owing to its better scalability over conventional forms of on-chip interconnect, improved performance, built-in fault tolerance, better load handling ability, modular topology to connect the processing elements, concurrent communication between several pairs of processing cores and improved parallelism [1], [2].

Scalable homogeneous NoC architecture comprises of several processing cores integrated on a single chip which are interconnected by a two dimensional mesh topology. High-speed routers, network interfaces (NI) and communication links are the basic components of regular tile-based NoC. The processing cores are attached to routers and the routers are connected by bi-directional links to exchange data between various processing elements (PE) in the form of packets. A packet is subdivided into flits (flow control units), which is the basic unit of data transfer in NoC. NoC network traffic is due to cache misses and coherence transactions. Each router has 5 bi-directional ports, North, South, East, West that are connected to neighbouring routers and a local port attached to corresponding PE [1], [3]. The conventional VCR based NoC design employ input buffers to improve throughput, network bandwidth utilization rate, load handling ability and thereby raise network performance. 2D mesh NoC architecture is widely used in TCMP due to its regular structure and scalability.

In this paper, we propose a novel design approach based on traditional VCR based NoC, where two layers of 2D mesh network are stacked on top of each other, without altering standard 5-port router microarchitecture. Comparison of our proposed design with a planar 2D mesh network and 2D torus network engaging equal number of routers show improved network saturation point and throughput with minimal average latency, power consumption and footprint while running at similar frequency as conventional 2D design.

The rest of this paper is organized as follows: In Section II gives an overview about the related work and section III describes motivation behind our design. Section IV present features of proposed design. Section V gives a description about experimental methodology followed. The results are discussed in Section VI and section VII concludes the paper.

2 Related Work

As 2D integrated circuits (IC) have limited floor planning choices, performance improvements occurring from NoC designs is restricted. The developments made in 3D IC technology has paved the way for routers used in 2D NoC topology to migrate towards 3D based topology. 2D NoC architectures use links made of global copper wires whereas 3D NoC comprises of both copper links and vertical TSV interconnects. 3D ICs can improve system performance manifold as they consists of several layers of active devices. Due to decrease in interconnect length, 3D ICs has improved performance, reduced power due to shorter interconnects, better packaging density and greater noise immunity [4], [5].

Stacked Mesh or Hybrid 3D NoC-Bus mesh structure is proposed in [6] which is a combination of packet-switched network and bus structure. 7-port symmetric 3D NoC router is replaced with 6-port hybrid NoC-Bus 3D switch to exploit small interlayer distance in a 3D IC for improving performance. Their approach uses addition arbiter for each pillar or vertical bus for improved integration between NoC and bus structures. The authors in [7] compare 2D mesh NoC structures and corresponding 3D structures by evaluating power consumption, speed and zero-load latency to indicate 3D NoC advantages over 2D NoC. They also developed an analytical model for zero-load latency of each network under consideration which takes into account the topology effects on 3D NoC performance. An exhaustive study of inter-strata communication architectures in 3D NoC is described in [8]. Several design options such as a hop-by-hop symmetric 3D design, a simple bus-based vertical connection and a 3D crossbar structure are explored in their work. They also propose an improved partially-connected 3D crossbar structure called DimDe, to deliver better performance and energy-delay product characteristics.

MIRA (Multi-layered on-chip Interconnect Router Architecture) proposed by [9] is a stacked 3D NoC structure which is optimized to minimize power consumption and the area requirements thereby enhancing performance and thermal behaviour. Feero et. al. [10] analyse performance of different 3D NoC architectures to exhibit their improved functionality in contrast to conventional 2D structures. They have introduced a novel architecture termed as ciliate 3D mesh which is basically a 3D mesh with several IP blocks per each switch. Each router is comprised of 7 ports, but their architecture has reduced overall bandwidth owing to multiple IP blocks per router and minimal connectivity when compared to a complete 3D mesh network. The authors in [11] report several application mapping and TSV placement strategies for 3D NoC systems. They also propose exact and heuristic techniques to address thermal-aware system design and test methods for both 2D and 3D NoC based architectures.

The effects of minimizing number of Through-Silicon-Via (TSV) to half and quarter on functionality of 3D NoC are assessed in [12]. Unbalanced distribution of 3D switches and random delays for different applications are main drawbacks of their architectures. Authors in [13] employ partition islands of switches to form areas for allocating same TSV pad for communication between interlayers, which are managed by serialization logic. But the average packet delay tends to exponentially rise with increase in number of switches per TSV bundle owing to serialization across TSV bundle. Contrary to 3D NoC employing vertical arbitration, a novel arbitration-free design is proposed for shared vertical links in [14]. Their proposed design has better performance in energy, throughput and latency when compared with symmetric 3D NoC with same area footprint.

3 Motivation

In a 2D NoC architecture, as the number of processing cores grows, there is significant increase in average packet latency whereas the communication quality

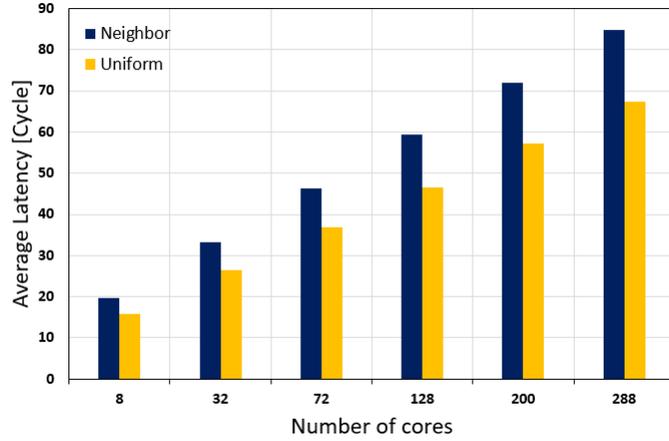


Fig. 1. Latency versus number of processing cores for different synthetic traffic patterns

decline due to rise in average number of hops between routers. Figure 1 clearly depicts the rise in flit latency with increasing number of cores for different synthetic traffic patterns in a planar 2D mesh NoC network. In contrast to 2D NoC network, 3D NoC designs will have minimal packet latency and smaller area footprint. But, 3D integration in NoC incur extra challenges as packets have to traverse along third dimension for which 3D router architectures have to be used. This increases arbitration, number of interconnections and ports in routers. The authors in [15], clearly indicate that crossbars used in 3D routers have increased power dissipation and area in contrast to 2D NoC routers.

The performance and area benefits of 2D NoC and 3D NoC architectures are put to use in our approach for enhancing functionality of 2D virtual channel routers. 5-port router architecture of conventional VCR based NoC is used in our design to build 2 identical layers of 8x8 meshes. The inter-layer communication is by vertical interconnections using TSVs made at edge routers through their unused ports thereby, improving routing efficiency.

4 Proposed Design

As mesh topology has a regular structure and routers are interconnected by short wires, we have taken mesh network for our proposed work. Figure 2 shows 8x16 2D mesh NoC and figure 3 details our new structure: 2L-2D. The proposed structure is formed by stacking 2 layers of 8x8 2D mesh NoC network. Both the layers are connected only through edge routers using vertical interconnect such as TSV. In a conventional 2D mesh NoC, all the 5-ports of routers, excluding the edge routers, are fully utilized in forming on-chip interconnect structure. But the edge routers have unused ports, which we have utilized in our design,

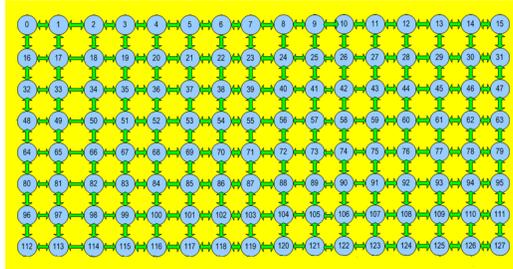


Fig. 2. 8x16 2D Mesh NoC

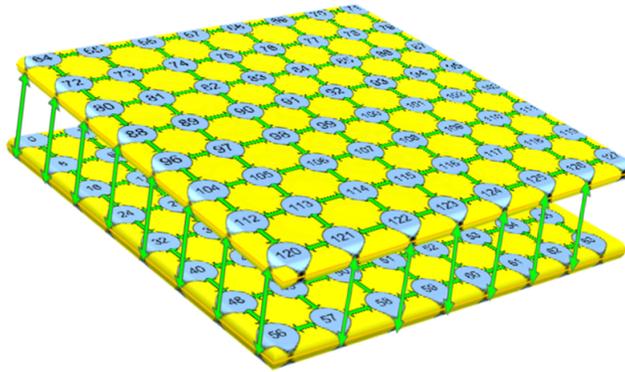


Fig. 3. 2L-2D 8x8 Mesh NoC

to make connection between two layers through vertical interconnection links. Thus instead of having a 7-port router structure as in a 3D mesh NoC network, routers in our proposed design can have the same 5-port router architecture as in a 2D mesh NoC.

The 5-port router used in this design is same as that of traditional virtual channel router as shown in figure 4. Conventional VC routers employ buffers at their input ports so that flits can stay in them till they attain a productive port. Complex buffer management circuitry is required in addition to buffers, consuming significant portion of on-chip power and occupying large footprint on the chip [16], [17]. The routing unit calculates required output port for the packet. Virtual Channel (VC) allocator unit does the task of reserving a VC in downstream router for each packet. Switch Allocation (SA) module performs arbitration to pick the winning packet when several packets require same downstream router. 5x5 crossbar is employed as switching fabric.

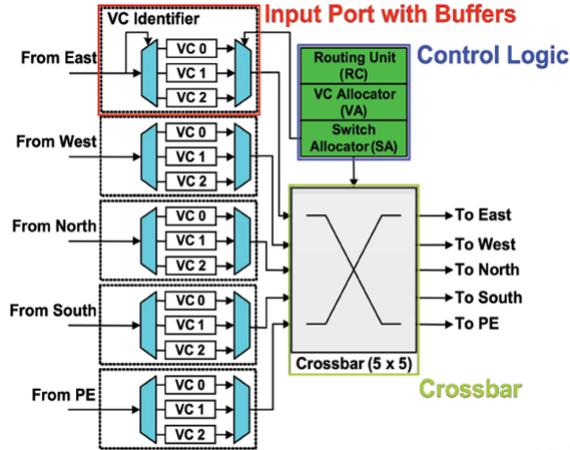


Fig. 4. Virtual Channel Router of [1]

Flit is routed to nearby router based on XY routing algorithm in 8x16 2D mesh and torus NoC whereas a modified XY routing algorithm is employed in our design for efficient routing. Simple static XY routing is used for intra-layer routing. When the source and destination routers are in different layers, flits advance to nearest edge router in X-dimension which gives minimal Manhattan distance across source and destination routers, as shown in algorithm aside. The above design approach can be implemented for bufferless mesh NoC as well.

5 Experimental Methodology

We employ Booksim 2.0 [18], an open source cycle accurate NoC simulator which models conventional VC based NoC router [1]. For our evaluations, we have used folded torus topology as in [1], to remove lengthy end-around link at the cost of doubling the span of other links. We then make necessary alterations to model our proposed design.

5.1 Synthetic Workload

The performance of 2L-2D is compared against VCR for mesh and torus topologies with size 8x16 (128 nodes) using standard synthetic workloads like uniform, tornado, bit-reverse, bit-complement, shuffle, neighbor and hotspot. After sufficient warm up time, average latency of flits and throughput readings are taken for all the traffic patterns by varying injection rates between zero and network saturation. Due to space limitations, results plotted for few traffic patterns only are shown.

Algorithm 1: 2L-2D Routing algorithm

Input : current_router, destination_router
Output: output port

```
if (current_router_layer == destination_router_layer) then
| XY routing algorithm
else
| if current router is edge router then
| | if current router is at first column then
| | | output port = west
| | else if current router is at last column then
| | | output port = east
| | else if current router is at first row then
| | | output port = north
| | else if current router is at last row then
| | | output port = south
| | end
| else
| | output port taken as east or west depending on shortest Manhattan
| | distance between source and destination router
| end
end
```

5.2 Real Workloads

We also evaluate the performance of our new design against baseline VCR of size 8x16 for both mesh and torus topologies using real application mixes such as multi-programmed SPEC CPU2006 [19] benchmark application suite and multi-threaded PARSEC benchmarks [20]. Multi2Sim [21] simulator is used to model a 128-core multiprocessor system where every processing core comprises an out-of-order x86 processing module with 64KB, 4-way set-associative private L1 cache and 512KB, 16-way set-associative shared distributed L2 cache. One of the benchmark applications from SPEC CPU2006 suite is run on each processing core. Depending on misses per kilo instructions (MPKI) values, we divide applications into Low, Medium and High MPKI. 7 multiprogrammed workload mixes are produced by combining the different applications from application suite. We also run multithreaded workloads on an equivalent setup with minor alterations to create sufficient traffic for our analysis. Network events produced by running these workloads are then fed as traffic to NoC simulator.

6 Results and Analysis

The performance of 2L-2D is compared against VCR based NoC for both mesh and torus topologies of size 8x16 (128 nodes). For VC router (VCR), we assume 16 VCs per input port for our analysis purpose. We employ deadlock free, deterministic, minimal path XY routing algorithm in VCR to analyse the performance improvement of our proposed design.

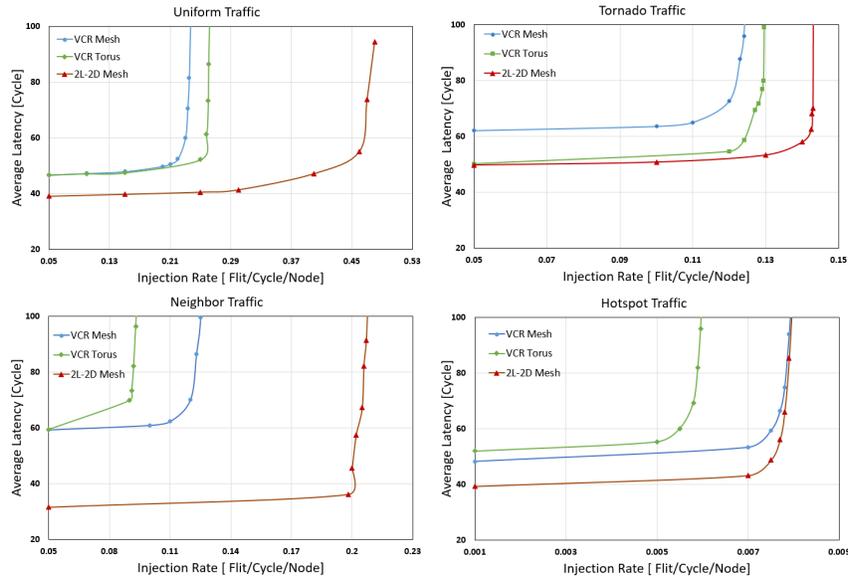


Fig. 5. Average latency comparison for different synthetic traffic patterns.

6.1 Effect on Average Latency

Flit latency is computed as the time instant at which flit is first created in network to the time instant when flit reaches destination core, including queuing time at source node. Wider and lower latency curve indicates that router performance is better. Average flit latency comparisons between 8x16 2D mesh VCR, 8x16 2D torus VCR and our proposed design under some of the typical synthetic traffic patterns are shown in figure 5. 2L-2D reduces average latency by 17%, 20% and 18% for uniform, tornado and hotspot traffic respectively, compared to VCR mesh. The average latency reduction of 2L-2D compared to VCR torus is 16%, 7% and 24% for the same traffic patterns. When injection rate nears saturation, there is exponential increase in average latency due to flooding of network with flits. Higher the saturation injection rate in a router better is its load handling capacity. It is quite evident from figure 5 that our proposed design has much improved saturation injection rate compared to VCR mesh and torus networks.

Figure 6 depicts comparison of average latency values as the number of processing cores is increased. On scaling number of cores, our 2L-2D design has better reduction in average latency over other two designs. The reduction in average latency of flits compared to VCR mesh and torus networks for multi-programmed and multithreaded workloads are shown in figure 7. There is significant decrease in average flit latency values for all the mixes using our design approach. As in multi-dimensional architectures, this reduced latency is due to

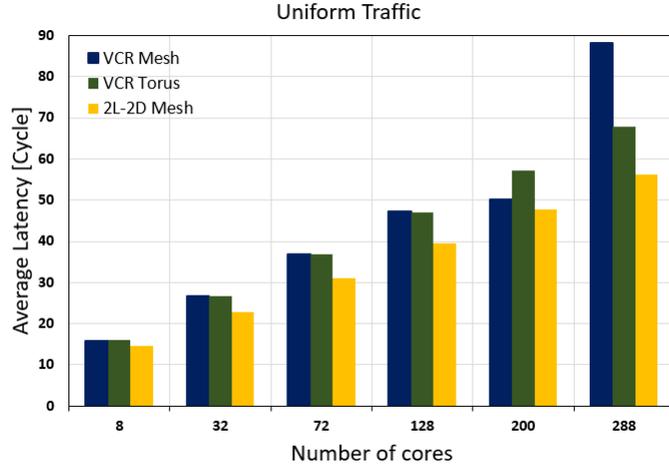


Fig. 6. Average latency versus number of processing cores

reduced hop count and more number of NoC links in our design where edge routers are interconnected through TSVs.

6.2 Area Overhead

In a NoC network, router overhead and wiring overhead contributes to total area overhead. Overall router area is based on total number of routers used in the network and area overhead per router, which is dependent on number of ports. 2D NoC architecture generally uses 5-port router architecture for mesh network whereas in 3D NoC designs, generally 7-port routers are utilized. Our proposed design has used same 5-port router architecture of 2D design thus, saving on router area overhead. There is negligible hardware router area overhead for our proposed work due to the modified routing logic which is overshadowed by remarkable reduction in average latency and throughput improvement.

For a 2D mesh NoC, wiring overhead includes overhead due to horizontal wirings only (8x16 2D mesh employs 232 horizontal links). For 3D NoC, in addition to area due to horizontal and vertical wirings, there is inter-layer via footprint. 2L-2D uses twice the number of horizontal links present in 8x8 mesh network (224 horizontal links) and some additional vertical wirings (28 vertical links) for connecting edge routers, thereby raising total wiring overhead. TSV has been utilized to connect inter-layer routers, which consumes some amount of silicon area and metal area. But our 2 layer design has reduced footprint as two 8x8 networks are stacked on top of each other.

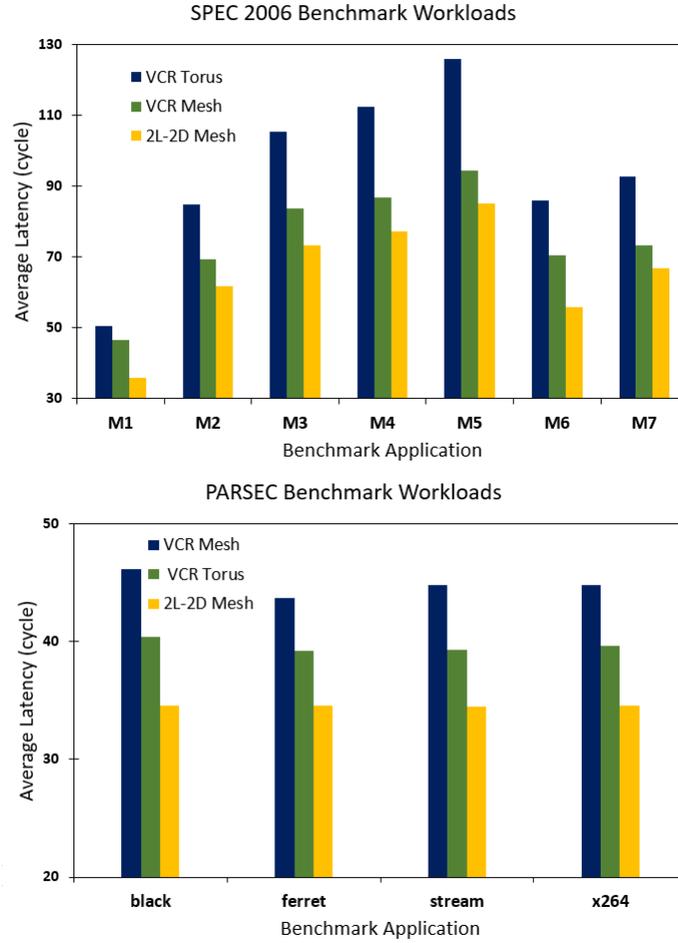


Fig. 7. Average latency for real applications

6.3 Effect on Throughput

Network throughput is calculated as number of flits ejected from network per router per cycle. For a multi-layer NoC network, throughput improvement depends on average hop count and number of physical links. Throughput delivered by an ideal NoC network will be same as the flit injection rate. As 2L-2D has more physical links and less average number of hops, our proposed design has better throughput indicating better quantity of sustainable traffic. Comparison with synthetic traffic patterns such as uniform, neighbor and bit-complement indicate throughput improvement of 2L-2D over other designs as shown in figure 8.

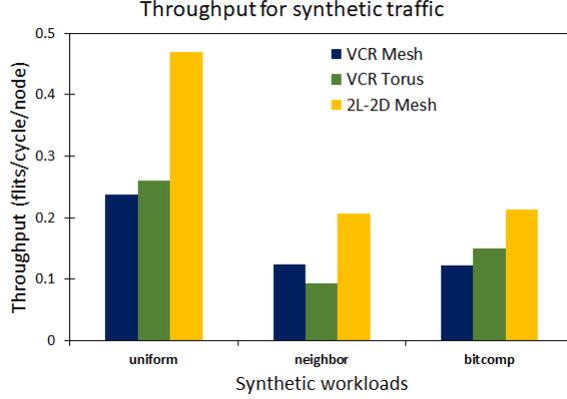


Fig. 8. Throughput comparison for different synthetic traffic patterns.

6.4 Effect on Router Pipeline Delay

Verilog HDL models of a router used in VCR based NoC and our proposed work is implemented and synthesized using Xilinx Vivado Design Suite-HLx to calculate pipeline latency of the router. Router delay is calculated as the time that a flit takes to traverse across router input port to router output port. The same functional units that are used in VCR are also employed in our design since both of them make use of 5-port router structure, only difference being in the routing logic used. Thus, our router pipeline frequency also remains the same as that of conventional VCR NoC router as similar functional units are employed in our design.

6.5 Effect on Dynamic Power Consumption across NoC links

Power dissipation in a NoC network rests on power dissipation across routers and inter-router wire links, which in turn rely on underlying, interconnect architecture. Power consumption in any network will rise as injection rate is increased since it determines amount of activity in routers and inter-router wires. Power dissipation of each flit per hop can be specified as

$$P_{flithop} = P_{router} + P_{link} \quad (1)$$

where P_{router} denotes power dissipation across each router and P_{link} is across inter-router links. Since identical routers are used for both 8x16 VCR and our design, power dissipation across router is going to be the same. Power dissipation of a flit which takes n number of hops can be given as

$$P_{flit} = \sum_{i=1}^n P_{flithop, i} \quad (2)$$

The average power dissipation per flit when N flits are transmitted is specified as

$$Pflitavg = \frac{1}{N} \sum_{j=1}^N Pflit \quad (3)$$

Thus average power dissipation of each flit depends on number of hops between the source and destination nodes.

The area and power is calculated and compared using Orion [22]. We presume 65nm technology at 1GHz operational frequency with one cycle inter-router link delay. Router area is same for all the designs as we have employed the same 5-port router microarchitecture in our design also. Dynamic power dissipation reduces by 23% for uniform traffic, 26% for tornado and 25% for hotspot traffic when compared with VCR mesh network. There is power dissipation reduction of 7%, 8% and 16% for uniform, tornado and hotspot traffic respectively when compared with torus network. Thus, our proposed design has significant reduction in dynamic power consumption across NoC links due to reduced average number of hops.

7 Conclusion

NoC has emerged as a promising solution to overcome scalability and bottleneck challenges faced by conventional SoC architectures. In this paper, we proposed 2L-2D, where two similar layers of 8x8 mesh NoC are placed on top of each other and minimal TSV based vertical interconnections are made through the unused ports of edge routers only. The performance and area advantages of 2D and 3D architectures are exploited in our approach using minimal number of TSVs and employing conventional 5-port VCR based NoC router. This design can be further extended to 3D NoC structure where multiple layers are stacked on top of each other and communication between layers are only through edge routers so that the same 5-port router architecture can be employed. The thermal issues associated with our proposed design can be evaluated as a future work.

References

1. W. Dally *et al.*, *Principles and Practices of Interconnection Networks*, Morgan Kaufmann, USA, 2004.
2. W. Dally, "Route packets, not wires: On-Chip interconnection networks", in *Design Automation Conference (DAC-01)*, pages 684-689, New York, ACM Press, June 2001. doi: 10.1109/DAC.2001.156225
3. W. Dally, "Virtual-channel flow control," *IEEE Transactions on Parallel and Distributed Systems*, vol. 3, no. 2, pp. 194-205, 1992. doi: 10.1109/71.127260
4. A. W. Topol *et al.*, "Three-Dimensional Integrated Circuits," in *IBM J. Research and Development*, Vol. 50, No. 4/5, 2006. doi: 10.1147/rd.504.0491
5. W. R. Davis *et al.*, "Demystifying 3D ICS: The Pros and Cons of Going Vertical," in *IEEE Design and Test of Computers*, Vol. 22, No. 6, pp. 498-510, 2005. doi: 10.1109/MDT.2005.136

6. F. Li *et al.*, "Design and management of 3D chip multiprocessors using network-in-memory," in *Proc. Int. Symp. Comput. Archit.*, pp. 130141, 2006. doi: 10.1109/ISCA.2006.18
7. V.F. Pavlidis *et al.*, "3-D Topologies for Networks-on-Chip," in *IEEE Trans. Very Large Scale Integration (VLSI '07)*, pp. 1081-1090, 2007. doi: 10.1109/TVLSI.2007.893649
8. J. Kim *et al.*, "A Novel Dimensionally-Decomposed Router for On-Chip Communication in 3D Architectures," in *Proc. of International Symposium on Computer Architecture*, pp. 138-149, 2007. doi: 10.1145/1273440.1250680
9. D. Park *et al.*, "MIRA: A Multi-layered on-chip Interconnect Router Architecture," in *Proc. of International Symposium on Computer Architecture*, pp. 251-261, 2008. doi: 10.1109/ISCA.2008.13
10. B. S. Feero *et al.*, "Networks-on-Chip in a Three-Dimensional Environment: A Performance Evaluation," in *IEEE Transactions on Computers*, pp. 32-45, 2009. doi: 10.1109/TC.2008.142
11. K. Manna *et al.*, "Thermal-aware Design and Test Techniques for Two- and Three-Dimensional Networks-on-Chip," in *2016 ISVLSI*, pp. 583-586, 2016. doi: 10.1109/ISVLSI.2016.76
12. T. Xu *et al.*, "A study of through silicon via impact to 3D network-on-chip design," in *Proc. Conf. Electron. Inf. Eng.*, pp. 333-337, 2010. doi: 10.1109/ICEIE.2010.5559865
13. Y. Wang *et al.*, "Economizing TSV resources in 3D Network-on-chip design," in *IEEE Trans. Very Large Scale Integration Syst.*, vol. 23, no. 3, pp. 493506, Mar. 2015. doi: 10.1109/TVLSI.2014.2311835
14. A. More *et al.*, "Vertical Arbitration-Free 3-D NoCs," in *IEEE Trans on Computer-Aided Design of Integrated Circuits and Systems*, vol. 37, no. 9, pp. 1853-1866, Sept. 2018. doi: 10.1109/TCAD.2017.2768415
15. M. O. Agyeman *et al.*, "Performance and Energy Aware Inhomogeneous 3D Networks-on-Chip Architecture Generation," in *IEEE Transactions on Parallel and Distributed Systems*, Vol.27, No.6, pp. 1756-1769, 2016. doi: 10.1109/TPDS.2015.2457444
16. Y. Hoskote *et al.*, "A 5-GHz mesh interconnect for a teraflops processor," *IEEE Micro*, vol. 27, no. 5, pp. 51-61, 2007. doi: 10.1109/MM.2007.4378783
17. M. B. Taylor *et al.*, "Evaluation of the raw microprocessor: An exposed wire-delay architecture for ILP and streams," in *ISCA*, 2004. doi: 10.1109/ISCA.2004.1310759
18. N. Jiang *et al.*, "A Detailed and Flexible Cycle-Accurate Network-on-Chip Simulator," in *IEEE International Symposium on Performance Analysis of Systems and Software*, 2013. doi: 10.1109/ISPASS.2013.6557149
19. "SPEC2006 CPU benchmark suite," <http://www.spec.org>.
20. C. Bienia *et al.*, "The parsec benchmark suite: characterization and architectural implications," in *PACT*, pp. 7281, 2008.
21. R. Ubal *et al.*, "Multi2sim: A simulation framework to evaluate multicore-multithreaded processors," in *SBAC-PAD*, pp. 62-68, 2007. doi: 10.1109/SBAC-PAD.2007.17
22. A. B. Kahng *et al.*, "ORION 2.0: A Fast and Accurate NoC Power and Area Model for Early-Stage Design Space Exploration," in *Design, Automation Test in Europe (DATE)*, pp. 423-428, 2009. doi: 10.1109/DATE.2009.5090700