Asymmetric Routing in 3D NoC using Interleaved Edge Routers

Rose George Kunthara
Rekha K James
Division of Electronics
School of Engineering, CUSAT
Cochin, India
{rosegeorgekunthara,rekhajames}@cusat.ac.in

Simi Zerine Sleeba
Department of Electronics & Communication Engineering
Viswajyothi College of Engineering and Technology, Muvattupuzha, India
simi.abie@gmail.com

John Jose
Department of Computer Science and Engineering
Indian Institute of Technology Guwahati, India
johnjose@iitg.ac.in

ABSTRACT

Network on Chip (NoC) concept has evolved as a standard design approach for integrating large number of processing cores within a single die. The performance improvements which occur due to topological optimizations and architectural enhancements of NoCs can be remarkably increased by the adoption of 3D IC fabrication technology. However, conventional 3D NoC architecture designs are significantly affected by router area and power dissipation issues when compared with traditional 2D NoCs. Also, 3D NoC utilizes considerable number of Through Silicon Vias (TSVs) which raises area overhead leading to minimal yield and wafer utilization. In this paper, we propose an asymmetric routing approach in bufferless 3D NoC using interleaved edge routers for enhancing NoC performance. Simulation results show that our proposed M-3D (Modified Three Dimensional) mesh design has better throughput, lower average flit latency and deflection rate compared to state-of-the-art bufferless networks, employing same number of routers.

CCS CONCEPTS

• **Computer systems organization → Multicore architectures**.

KEYWORDS

Bufferless routing, Flit latency, Throughput, Deflection routing

ACM Reference Format:
Rose George Kunthara, Rekha K James, Simi Zerine Sleeba, and John Jose. 2019. Asymmetric Routing in 3D NoC using Interleaved Edge Routers. In *12th International Workshop on Network on Chip Architectures (NoCArc '19), October 13, 2019, Columbus, OH, USA.* ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3356045.3360720

1 INTRODUCTION

The exponential rise in transistor integration density has resulted in a major overhaul from off-chip centric design to on-chip centric design approach. This has led to the advent of Tiled Chip Multi Processors (TCMP) where multiple processing cores are integrated on a single chip. According to International Technology Roadmap for Semiconductors (ITRS), with technology scaling from 7nm to 1nm in 2028, thousands of intellectual property (IP) cores can be integrated on a single chip [1]. NoC has emerged as the widely accepted communication backbone replacing conventional bus based and point-to-point intercommunication architectures, owing to its better scalability, reliability, modular topology, load handling ability and concurrent communication among several cores [2][3].

In NoC, conventional input buffered Virtual Channel Router (VCR) has better load handling capability and network performance as flit remains in buffer until it acquires a productive output port [2][4]. In addition to complex buffer management circuitry, buffers consume considerable fraction of on-chip area and power [5][6]. Experimental results show that for low to medium injection rates, alternate design options such as bufferless NoC routers outperform VCR based NoC design [7]. Bufferless routers employ deflection routing where flits which do not get desired output port will be assigned to available output ports. This can sometimes lead to increased deflection rate and latency. Two major bufferless routers based on deflection routing are BLESS [7] and CHIPPER [8].

We propose a hybrid design approach by incorporating 2D CHIPPER design and 3D NoC using TSV. Several layers of 2D mesh network are stacked using 3D integration. Our proposed M-3D design employs same 5-port router architecture of CHIPPER where interlayer communication is through TSV interconnections made only at edge routers. On comparing with 2D planar mesh and 3D mesh, our design shows better network performance with minimal footprint and router overhead while operating at the same frequency as 2D CHIPPER.

The remainder of this paper is structured as follows: Section II gives an overview of the related work and Section III details motivation for our proposed design. The proposed M-3D design is described in Section IV and the experimental methodology is discussed in Section V. Results and analysis is given in Section VI. Finally Section VII concludes the paper.

2 RELATED WORK

In NoC design, the restricted floor planning choices associated with 2D integrated circuits (IC) pose a performance bottleneck. As number of processing cores grows with shrinking geometries, global wire delay problem can be overcome by extending 2D NoC to 3D NoC systems by stacking the cores into various layers. 3D NoC can enhance the system performance manifold due to better noise immunity, packaging density and reduced power consumption because of short interconnect wires [9][10]. In 3D IC, TSVs are generally used for interlayer communication owing to their superior performance. 3D NoC employs 3D routers, which are natural extension of 2D routers with additional ports in z-dimension.

NoCArc '19, October 13, 2019, Columbus, OH, USA

ACM ISBN 978-1-4503-6949-7/19/10…$15.00
https://doi.org/10.1145/3356045.3360720

Pavlidis et al. has compared 2D mesh networks with their 3D equivalents to present advantages of 3D NoC by evaluating power consumption and zero-load latency of each network [11]. A class of 3D topologies termed as Xbar-connected network-on-tiers (XNoTs) containing several network layers that are connected through crossbar switches are detailed by Matsutani et al. to make optimum use of short delay and high density of inter-wafer links [12]. Various forms of XNoTs based topologies are evaluated in terms of different performance metrics to show their superior throughput even though it employs large vertical switches thereby, degrading its power efficiency. MIRA, a 3D stacked NoC router architecture employs multiple layers and optimized to reduce power dissipation and overall area requirements, while assuming that the processing cores are also designed in 3D [13].

Xu et al. assess the effect of minimizing number of TSVs to half and quarter on the functionality and performance of 3D NoC system [14]. Unbalanced distribution of 3D switches and variable delays for various applications are major drawbacks of their proposed architectures. Wang et al. employ partition islands of switches to create areas for assigning same TSV pad for interlayer communication that are managed by serialization logic [15]. Due to the serialization across TSV bundle, average packet delay tends to increase exponentially with rise in number of switches per TSV bundle.

3DMAX, a maximally adaptive routing technique for fully connected 3D mesh network realises effective on-chip communication while eliminating routing tables or VCs [16]. Since all routers are not vertically connected in 3D NoC due to high manufacturing cost of TSV, it provides partial connectivity along the vertical dimension. Vahdatpanah et al. proposes an efficient routing algorithm to evenly distribute traffic in partially connected 3D NoCs [17].

TSV serialization in 3D NoC poses an obstruction to performance as it reduces latency and available bandwidth of TSV links. Lee et al. proposes a deflection routing scheme that allows full TSV link utilization to minimize latency [18]. Their work uses TSV ejection/injection mechanism to overcome performance bottleneck at high traffic. Feng et al. proposes 3DPERM, a single cycle deflection router that has decreased area and power overhead compared to a single-cycle 3D CHIPPER [19]. 3DBUFFBLESS is another asymmetrical 3D NoC router that is bufferless in x and y dimensions whereas buffered in z-dimension [20]. Their proposed router improves routing efficiency with minimal area and power dissipation by combining advantages of bufferless and buffered router architectures.

3 MOTIVATION

Increasing number of processing cores in a 2D mesh topology leads to increase in latency, power dissipation, chip area and floor plan. We perform simulations on 2D mesh NoC using CHIPPER for different synthetic traffic patterns. Figure 1 shows average flit latency with rise in number of processing cores for uniform, tornado and neighbor traffic patterns at zero load. We can observe that, as network size increases, the average flit latency also increases even if at low network load. Since the large networks are spread across the 2D plane, this leads to larger power dissipation and poor performance.

To circumvent these limitations, several active NoC layers are stacked using vertical interconnects to form 3D NoC. TSV based

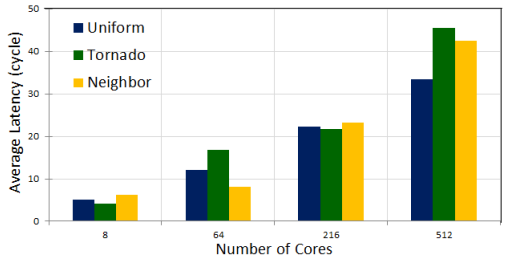

Figure 1: Latency versus number of processing cores for different synthetic traffic patterns at zero load

vertical interconnect is the most popularly used approach in 3D NoC for providing fast inter-layer communication due to its small form factor and better performance. But with increase in number of processing cores and subsequent rise in number of layers in 3D structure, more number of TSVs are to be used. This increases manufacturing cost, area overhead due to TSVs leading to reduced wafer utilization and yield in addition to misalignment and physical stress issues that affect the design of 3D NoC architectures [21].

7-port router architectures are mostly used in 3D NoC designs as the packets are to be routed in all the three dimensions. This increases the number of interconnections, arbitration and ports in the routers thereby increasing area overhead and power dissipation of routers. Generally, 3D NoC architectures have small area footprint and lower average latency compared to 2D NoC designs. But bufferless 3D mesh design, which uses golden flit scheme employed in CHIPPER, has higher average latency as the non-golden flits are arbitrated pseudo-randomly leading to more number of deflections.

The area and performance advantages of 2D NoC and 3D NoC designs are exploited in our proposed M-3D mesh approach for improving the scalability and performance of NoC. The 5-port router structure of CHIPPER is utilized in this design to create a 3D NoC architecture by a novel arrangement of vertical interconnects. For inter-layer communication, TSV based vertical interconnections are made only at edge routers by using their unused ports thereby, enhancing routing efficiency with minimal number of vertical interconnections and router hardware.

4 PROPOSED DESIGN

The topology, router microarchitecture and routing algorithm have a significant role in the performance of an NoC. Figure 2(a) shows a 2D 8x8 mesh NoC. 3D mesh is a natural extension of the popular planar 2D mesh structure with extra vertical ports for communication between adjacent layers. Figure 2(b) depicts a 3D mesh NoC that uses 7-port router architecture. Four of the router paths connect to the neighbouring routers in North, South, East and West direction, one port connect to the local processing core and the remaining two connect to routers in upper and lower layers.

Our proposed design uses mesh topology owing to its regularity, scalability and short interconnection wires. Figure 2(c) shows our proposed M-3D mesh NoC with 64 routers interconnected using a 4x4x4 mesh topology. Here adjacent layers are connected using interleaved edge routers. In a traditional 5-port 2D mesh NoC, each of the edge routers have one unused port. In M-3D mesh, we

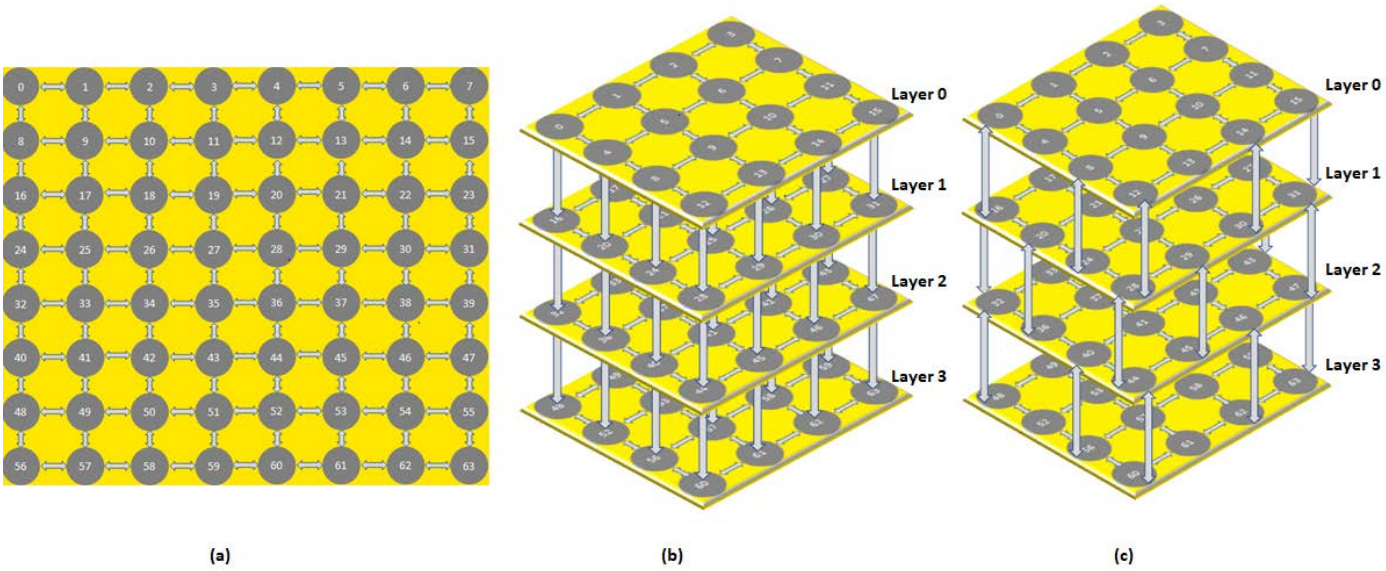

Figure 2: Mesh topology based NoC architectures. (a) 2D mesh (b) 3D mesh (c) Proposed M-3D mesh

utilize these ports to form interconnections between adjacent layers using TSV based vertical links. The interleaved connection between layers is illustrated using an example. Consider Layer 2 in Figure 2(c) where the 4 routers at one of the edges numbered as 44, 45, 46 and 47. The unused south port of Router 46 is interconnected to Router 62 in Layer 3; whereas Router 45 is connected to Router 29 in Layer 1. However corner routers of Layer 1 and Layer 2 are connected to both upper and lower layers. This type of interleaved interconnection is followed in routers at the other edges also.

M-3D mesh uses a two cycle bufferless router employing 5-port microarchitecture as in CHIPPER [8]. At every clock cycle, flits from adjacent routers arrive at input ports. The ejection and injection units constitute first stage of pipeline followed by a 2-stage permutation network in the second stage. Golden flit scheme ensures livelock avoidance. The golden flit, which is always treated as highest priority flit, will be assigned the desired output port whereas the remaining flits undergo random port assignment.

As flits are not buffered inside the router, deflection routing algorithms are deadlock free. Our proposed M-3D mesh design follows asymmetric routing as described in Algorithm 1. For intra-layer routing, the simple deterministic, static XY routing algorithm is followed. If the current router and destination routers are in separate layers, then flits are routed to nearest edge router which is connected to corresponding layer with shortest Manhattan distance. Thus, asymmetry exists in the routing path and in router interconnections due to interleaving which is adopted to distribute the connections evenly between all layers.

5 EXPERIMENTAL METHODOLOGY

We use Booksim 2.0, an open source cycle accurate NoC simulator that models traditional VC based NoC router [22]. Requisite modifications are done to prototype a two-cycle bufferless deflection router CHIPPER [8]. Necessary information is attached to every flit under consideration so as to promote independent routing of

Algorithm 1: Routing algorithm for proposed M-3D mesh

```
Input  : current_router, destination_router
Output : output port

if (current_router_layer == destination_router_layer) then
    XY routing algorithm
else if (current_router_layer is outer or odd) then
    if current router is at first column and even row then
        output port = west
    else if current router is at last column and odd row then
        output port = east
    else if current router is at first row and odd column then
        output port = north
    else if current router is at last row and even column then
        output port = south
    else
        output port taken as east or west depending on
        destination column
    end
else
    if current router is at first column and odd row then
        output port = west
    else if current router is at last column and even row then
        output port = east
    else if current router is at first row and even column then
        output port = north
    else if current router is at last row and odd column then
        output port = south
    else
        output port taken as east or west depending on
        destination column
    end
end
```

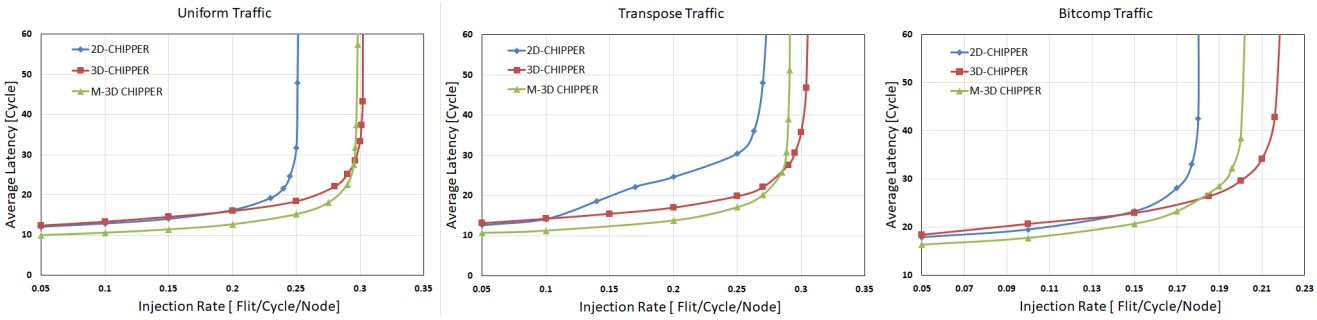

Figure 3: Average flit latency comparison for different synthetic traffic patterns

Table 1: Classification of benchmark applications

Percentage Miss Rate	Benchmark applications
Low: MPKI < 5	calculix, gobmk, gromacs, h264ref
Medium: $5 \leq MPKI \leq 25$	bwaves, bzip2, gamess, gcc
High: MPKI > 25	hmmer, lbm, leslie3d, mcf

Table 2: Percentage combination of different network injection intensity applications in various benchmark mixes

Benchmark Mix	M1	M2	M3	M4	M5	M6	M7
% of Low	100	0	0	50	0	50	31
% of Medium	0	100	0	0	50	50	31
% of High	0	0	100	50	50	0	38

all flits inside a packet, as is the usual norm in bufferless deflection routers. Out-of-order flit delivery is handled by requisite reassembly mechanism. 140-bits wide flit channel is constituted by 128-bits data field and 12-bits control field. We make necessary modifications on this baseline bufferless deflection router simulator to prototype 3D CHIPPER as well as M-3D CHIPPER and conduct experimental analysis.

5.1 Synthetic Traffic

We evaluate performance of M-3D CHIPPER against 2D CHIPPER and 3D CHIPPER with 64 routers using standard synthetic traffic patterns like uniform, transpose, tornado, bit-complement, bit-reverse, neighbor and shuffle. Average flit latency, deflection rate and throughput are computed for all traffic patterns after sufficient warm up time by varying injection rate from zero to network saturation point.

5.2 Real Traffic

The performance of M-3D CHIPPER is compared against baseline 2D CHIPPER and 3D CHIPPER employing same number of nodes (64 nodes), for real application mixes consisting of multiprogrammed SPEC CPU2006 benchmark applications [24]. Gem5 simulator is used for modelling a 64-core multiprocessor system with each processing core consisting of an out-of-order x86 processing module with 4-way set associative, 64KB private L1 cache and 16-way set associative, 512KB shared distributed L2 cache [23]. Each processing core runs one of the benchmark applications from SPEC CPU2006 application suite. The benchmark applications are categorized into different network injection intensity classes based on values of misses per kilo instructions (MPKI) as Low, Medium and High MPKI. Details of benchmark classification are shown in Table 1. We produce 7 multiprogrammed workload mixes depending on network injection intensity proportion of component benchmarks as shown in Table 2. To simulate network operations, cache miss requests and reply packets generated from gem5 is fed to Booksim and statistics are collected.

6 EXPERIMENTAL ANALYSIS

Our proposed design is compared against 2D CHIPPER and 3D CHIPPER in terms of average flit latency, average deflection rate and throughput. We conduct experiments on 8x8 mesh with 2D CHIPPER and XY routing, 4x4x4 mesh with 3D CHIPPER and XYZ routing, and 4x4x4 mesh with M-3D CHIPPER employing routing algorithm proposed in Algorithm 1.

6.1 Effect on Average Flit Latency

Flit latency refers to time elapsed between creation of flit in the network at source node and reception of flit at destination node. Figure 3 shows average flit latency comparisons using synthetic traffic patterns. Broader and lower flit latency curve is desirable for better router performance. Generally a reduction in latency is expected in 3D topology due to lower number of hops. But our latency graphs depict higher latency values for 3D CHIPPER due to random port assignment of non-golden flits which lead to more deflections and hence larger latency. From the plots it is quite evident that our proposed design approach has lower flit latency over all the traffic patterns.

An exponential rise in average latency occurs as the flit injection rate approaches saturation load. Larger saturation injection rate of a router indicates its superior load handling ability. M-3D CHIPPER improves the network saturation point than 2D CHIPPER and is almost at par with 3D CHIPPER. As the number of routers scales up, M-3D CHIPPER has significant latency reduction when compared to other designs as depicted in Figure 5. Figure 6 shows average flit latency for real applications. We can see that M1 experiences the minimum latency and M3 have maximum latency as expected. M3 has higher load due to the presence of heavy applications. Across

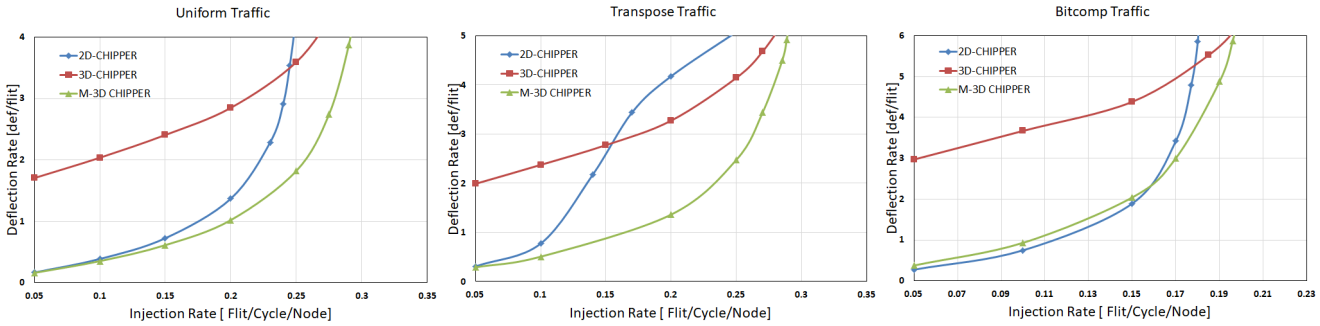

Figure 4: Average flit deflection rate comparison for different synthetic traffic patterns.

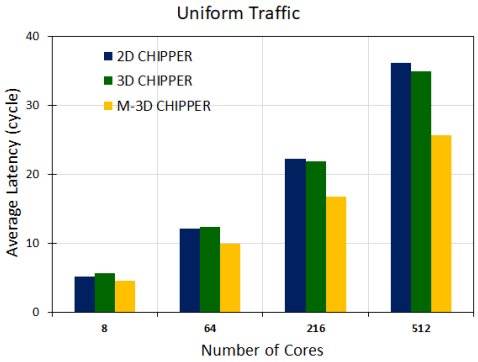

Figure 5: Average flit latency comparison under uniform traffic for different network sizes

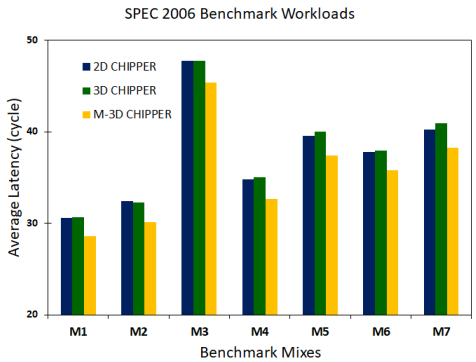

Figure 6: Average latency comparison for real applications

all mixes, M-3D CHIPPER experiences lower latency compared to other two designs.

6.2 Effect on Deflection Rate

Deflection rate is the average number of deflections that occur for each injected flit. A lower deflection rate implies reduced dynamic power across NoC links as a result of minimal network activity. With the increase in injection rate, deflection rate will also rise owing to higher port contentions. Figure 4 shows the deflection rate for synthetic traffic patterns. M-3D CHIPPER has the lowest deflection rate and this is more pronounced at higher injection rates. Figure 7 shows the deflection rate for SPEC CPU 2006 benchmark workloads. The deflection rate of 3D CHIPPER is omitted from graph as it is very high due to random port allocation of non-golden flits, compared to 2D CHIPPER and M-3D CHIPPER. It is obvious that for all the mixes, proposed M-3D CHIPPER has significant reduction in deflection rate than 2D CHIPPER.

6.3 Effect on Throughput

Throughput is calculated as number of ejected flits from the network per node per cycle. It is a measure of peak data rate which the system is able to handle. Average throughput denotes the rate at which traffic flows across the interconnection fabric. For an ideal bufferless network, throughput and flit injection rate will be the same. Multi-layer NoC network is expected to have better throughput due to higher number of physical (both horizontal and vertical) links. Figure 8 indicates that M-3D CHIPPER has better throughput compared to 2D CHIPPER indicating better sustainable traffic for various synthetic traffic. Throughput of M-3D CHIPPER is slightly lesser than that of 3D CHIPPER due to smaller number of TSVs used in it and their uneven distribution to form vertical interconnections.

6.4 Effect on Router Pipeline Latency, Power and Area

We implement Verilog HDL models of 2D CHIPPER, 3D CHIPPER and M-3D CHIPPER and synthesize using Synopsys Design Compiler with 90nm cell library to compute the router pipeline latency, power and area overhead. Overall time taken by a flit to move from input port to output port constitutes the router delay. Except for the routing logic, M-3D CHIPPER uses similar functional units as employed in 2D CHIPPER. As the second stage of 2D CHIPPER determines the critical path, proposed M-3D CHIPPER will have same operational frequency as that of 2D CHIPPER. There is a 52% reduction in critical path latency and 32% reduction in power consumption when compared to 3D CHIPPER.

Router area and wiring overhead constitutes the overall area overhead in any NoC network. Router area depends on internal microarchitecture and input ports. The proposed M-3D CHIPPER employs 5-port router architecture similar to that of 2D CHIPPER thereby, achieving 30% savings on router area when compared to 3D

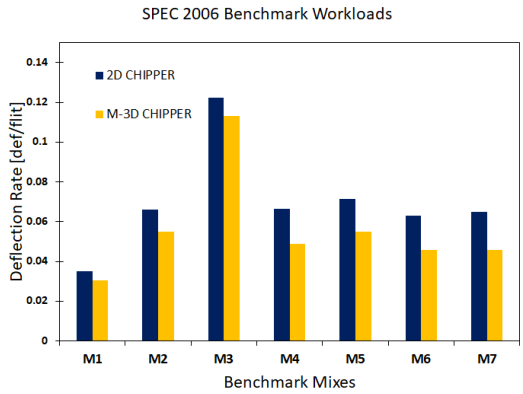

Figure 7: Average deflection rate comparison for real applications

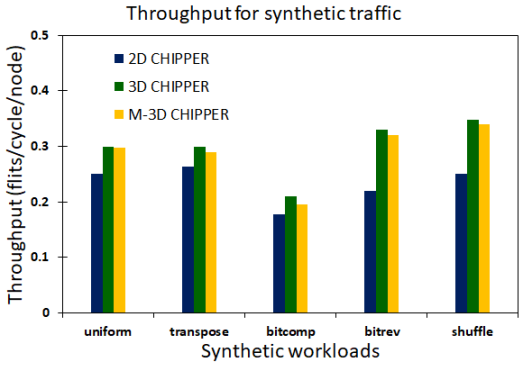

Figure 8: Throughput comparison for different synthetic traffic patterns.

CHIPPER, which uses 7-port router structure. Due to the modified routing logic used in our proposed design, there is negligible router area overhead of 2.3% and power dissipation of 2.6% compared to 2D CHIPPER which is greatly masked by the significant improvement in throughput and considerable reduction in average deflection rate and latency.

The wiring overhead of a 2D mesh NoC contains overhead owing to horizontal links only (8x8 2D mesh uses 112 horizontal links). In a conventional 4x4x4 3D mesh NoC, there are 96 horizontal links and 48 vertical links. Our proposed design uses the same number of horizontal links (96 links) as that of 3D network but reduced number of vertical links (24 links), as vertical interconnections are made only at edge routers. As TSVs consume significant metal area and silicon area, M-3D mesh has better area savings as we have used minimal number of vertical interconnections using TSVs. Overall our proposed M-3D CHIPPER has the same chip footprint as that of 3D mesh but with 50% less vertical links.

7 CONCLUSION

With massive rise in system integration, low cost solutions are required to deal with area, power and performance bottlenecks occurring in on-chip networks. 3D NoC has evolved as a promising approach to improve system performance by providing higher integration density and scalable communication platform for TCMP. In this paper, we have proposed M-3D mesh, a 3D NoC structure with interleaved edge routers using asymmetric routing scheme for improving NoC performance by utilizing minimal number of TSV based vertical interconnections and employing 5-port router architecture. Simulation results indicate that M-3D CHIPPER has lower latency and can sustain more network traffic compared to designs under consideration with minimum hardware modification.

REFERENCES

[1] Semiconductor Industry Association, The International Technology Roadmap for Semiconductors, ITRS 2015 Edition.
[2] W. Dally and B. Towles, Principles and practices of interconnection networks, San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 2003.
[3] W. Dally and B. Towles, Route packets, not wires: on-chip interconnection networks, In DAC, Pages. 684-689, 2001.
[4] W. Dally, Virtual-channel flow control, IEEE-TPDS, Vol. 3, Number.2, Pages 194-205, March 1992.
[5] Yatin Hoskote, Sriram Vangal, Arvind Singh, Nitin Borkar and Shekhar Borkar, A 5-GHz mesh interconnect for a teraflops processor, IEEE Micro, Vol. 27, Number. 5, Pages 51-61, 2007.
[6] M. B. Taylor et al., Evaluation of the raw microprocessor: An exposedwire-delay architecture for ILP and streams, In ISCA, 2004.
[7] Thomas Moscibroda and Onur Mutlu, A case for bufferless routing in on-chip networks, In ISCA, Pages 196-207, 2009.
[8] Chris Fallin, Chris Craik and Onur Mutlu, CHIPPER: A low complexity bufferless deflection router, In HPCA, Pages 144-155, 2011.
[9] A. W. Topol et al., Three-Dimensional Integrated Circuits, IBM Journal of Research and Development, Vol. 50, Number. 4/5, July 2006.
[10] W. R. Davis et al., Demystifying 3D ICS: The Pros and Cons of Going Vertical, IEEE Design and Test of Computers, Vol. 22, Number. 6, Pages 498-510, 2005.
[11] Vasilis F. Pavlidis and Eby G. Friedman, 3-D Topologies for Networks-on-Chip, IEEE-TVLSI, Vol. 15, Number. 10, Pages 1081-1090, October 2007.
[12] Hiroki Matsutani, Michihiro Koibuchi and Hideharu Amano, Tightly-Coupled Multi-Layer Topologies for 3D NoCs, In ICPP, 2007.
[13] D. Park et al., MIRA: A multi-layered on-chip interconnect router architecture, In ISCA, Pages. 251-261, 2008.
[14] T. Xu, P. Liljeberg, and H. Tenhunen, A study of Through Silicon Via impact to 3D Network-on-Chip design, In ICEIE, Pages. 333-337, 2010.
[15] Y. Wang et al., Economizing TSV Resources in 3-D Network-on-Chip Design, IEEE-TVLSI, Vol. 23, Number. 3, Pages. 493âĂŞ506, March 2015.
[16] Poona Bahrebar and Dirk Stroobandt, 3D MAX: A Maximally Adaptive Routing Method for VC-less 3D Mesh-based Networks-on-Chip, In NoCArc, 2018.
[17] F. Vahdatpanah, M. Elahi, S. Kashi, E. Taheri and A. Patooghy, 3DEP: A Efficient Routing Algorithm to Evenly Distribute Traffic Over 3D Network-on-Chips, In PDP2019, February 2019.
[18] Jinho Lee, Dongwoo Lee, Sunwook Kim and Kiyoung Choi, Deflection routing in 3D Network-on-Chip with TSV serialization, In ASP-DAC, 2013.
[19] C. Feng, Z. Lu, A. Jantsch, and M. Zhang, A 1-Cycle 1.25GHz Bufferless Router for 3D Network-on-Chip, IEICE Transactions on Information and Systems, Volume E95.D, Issue 5, Pages. 1519-1522, 2012.
[20] K. Tatas, S. Savva and C. Kyriacou, 3DBUFFBLESS: A Novel Buffered-Bufferless Hybrid Router for 3D Networks-on-Chip, In PATMOS, 2017.
[21] Michael Opoku Agyeman, Ali Ahmadinia and Nader Bagherzadeh, Performance and Energy Aware Inhomogeneous 3D Networks-on-Chip Architecture Generation, IEEE-TPDS, Vol.27, Number.6, Pages. 1756-1769, 2016.
[22] Nan Jiang, Daniel U. Becker, George Michelogiannakis, James Balfour, Brian Towles, John Kim and William J. Dally, A detailed and flexible cycle-accurate Network-on-Chip simulator, In ISPASS, Pages 86-96, 2013.
[23] N. Binkert et al., The gem5 simulator, SIGARCH Computer Architecture News, Vol. 39, Number. 2, Pages. 1-7, 2011.
[24] SPEC2006 CPU benchmark suite, http://www.spec.org.