Cluster-based Topologies for 3D Stacked Architectures

Masoud Daneshtalab, Masoumeh Ebrahimi, Pasi Liljeberg, Juha Plosila, Hannu Tenhunen
Department of Information Technology, University of Turku
{masebr,masdan,pakrli,juplos,hanten}@utu.fi

ABSTRACT
As Three Dimensional Integrated Circuits (3D ICs) have been emerging as a viable candidate to achieve better performance and package, combining the benefits of 3D IC and Network-on-Chip (NoC) schemes provides a significant performance gain for 3D architectures. Through-Silicon-Via (TSV), employed for inter-layer connectivity (vertical channel)in 3D ICs, reduces wafer utilization and yield which impact design of 3D architectures using a large number of TSVs. In this paper, we propose two novel stacked topologies for 3D architectures to reduce the area overhead of TSVs and power dissipation on each layer with minimal performance penalty. The presented schemes benefit of clustering the mesh topology in order to mitigate TSV footprint on each stacked layer.

Categories and Subject Descriptors
C.2.1 [**COMPUTER-COMMUNICATION NETWORKS**]: Network Architecture and Design - Packet-switching networks.

General Terms
Algorithms, Performance, Design.

1. INTRODUCTION
The Three Dimensional (3D) integration has emerged as a potent solution to address the design complexity of MPSoC in 2D Integration Circuits (IC). 3D NoC topologies not only create scalable networks to address communication requirements in 3D ICs [1]-[3] but also are a crucial factor of 3D chips in terms of performance, cost, and energy consumption [1]. Various on-chip network topologies have been studied for 3D NoCs [1]-[4][6]. Mesh-based structures are popularly used in 3D systems, because their grid-based regular structure is intuitively considered to be matched to the 2D VLSI layout for each stack layer [1][2][3]. Nevertheless, if the number of IP-cores and memories increases in each layer, more TSVs are necessitated to handle the inter-layer communication. Inasmuch as each TSV employs a pad for bonding, the area footprint of TSVs in each layer is augmented significantly [3][6].

In this paper, we propose two novel stacked mesh topologies to reduce the area overhead of TSVs and power dissipation with minimal performance penalty. The proposed stacked mesh topologies, named Clustered Mesh Inter-layer Topology (CMIT) and Concentrated Inter-layer Topology (CIT), benefit of clustering the mesh topology for each layer. Each cluster of the presented topologies has its dedicated vertical channel, composed of a set of TSVs. CMIT and CIT preserve the advantages of the clustered mesh topology and mitigates both power density and TSVs area footprint on each layer. The rest of the paper is organized as follows. In Section 2, the background and related work are discussed while the proposed topologies are presented in Section 3. The experimental results are discussed in Section 4 with the summary and conclusion given in the last section.

2. BACKGROUND AND RELATED WORK
3D-Symmetric NoC and 3D NoC-Bus Hybrid (stacked mesh) structures are popularly used in 3D systems, because their grid-based regular structure is intuitively considered to be matched to the 2D VLSI layout for each stack layer [1][2][3]. The 3D-symmetric NoC structure is an extension of 2D mesh by adding two additional physical ports to each baseline-router (one for up and one for down) in the popular 2D mesh-based system[1][3]. Adding two additional ports requires larger crossbar incurring significant area and power overhead and increases the blocking probability inside the router. Since TSVs are relatively shorter and wider than that of intra-layer interconnects, they have low resistance and can support higher signalling speeds [3][6]. As router latencies may dominate the throughput of fast vertical interconnects, this has led the researchers to propose the 3D NoC-Bus Hybrid structure using a bus with a centralized arbiter for each vertical channel, which allows single hop latency for packets between any number of layers [1]-[3]. On-chip routers in this structure have at most 6 ports, one to the IP-core, one to the bus, and four for cardinal directions. According to [3] the 3D hybrid structure was observed to be more efficient than the 3D symmetric for the vertical interconnection as long as the number of device layers was less than 9. It can also exploits the dynamic Time Division Multiple Access (dTDMA) bus [3] with a centralized arbiter for the vertical communication link. Thus, moving from one layer to any of the other layers takes only one hop and, contention issues in the bus limit the attainable performance gains [1].

3. CLUSTER-BASED ARCHITECTURES
Since each TSV requires a pad (around 5μ×5μ) with pitch of around 8μ for bonding to a wafer, the area overhead of TSVs impose constraints on the number of TSVs [3][6]. In order to reduce vertical channels, we present two novel topologies. Although both of the presented topologies can be formed as either the 3D-symmetric or 3D-hybrid structure, we describe these topologies based on the 3D-hybrid scheme which is more efficient than the 3D-structure [3][6]. Each vertical channel is composed of two unidirectional channels in opposite directions to propagate the inter-layer data.

3.1 CIT (Concentrated Inter-layer Topology)
Unlike the mesh topology where each IP-core is connected to a router, it forms a scalable architecture by sharing a router between multiple nodes (IP-cores and memories). CIT reduces the number of routers decreasing the number of vertical channels and hop counts. A 3×3 CIT with 36 nodes is shown in Fig. 1(a), where four nodes are grouped into a cluster, thereby 9 clusters are formed in

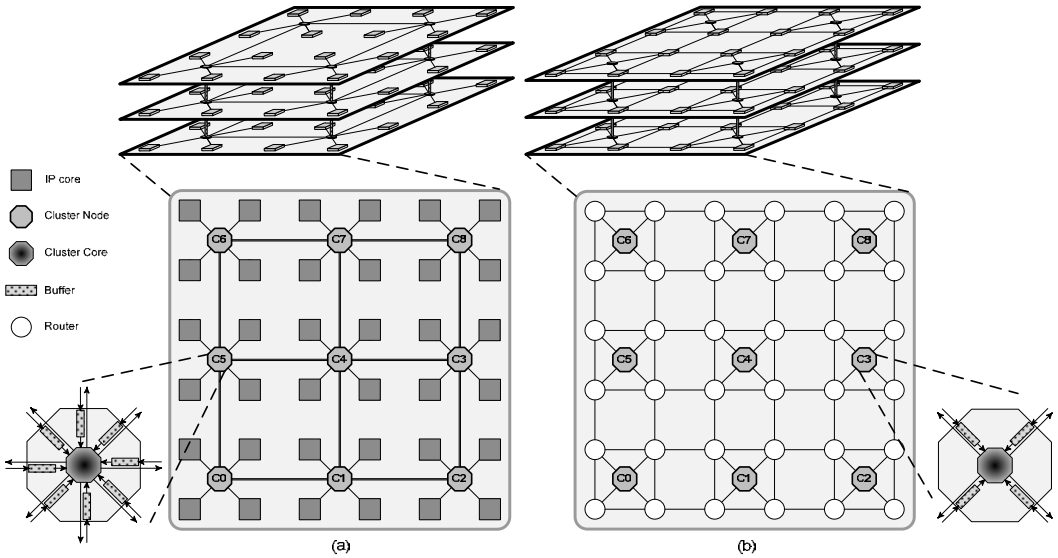

Fig. 1. Clustering approaches: (a) CIT and (b) CMIT.

the network. Each cluster has a router with at most 9 ports (10 ports for 3D-symmetric structure): one connected to the bus, four to IP-cores/memories and the other four ports to neighboring routers. Communication channels in CIT can be classified as intra-layer channels (horizontal channels) and inter-layer channels (vertical channels). As illustrated in Fig. 1, the inter-layer communication is achieved by the cluster nodes. Each cluster node has a cluster core to establish the vertical connection via an interface to the vertical bus. Due to the fact that each CIT router has larger number of input ports than the symmetric and hybrid routers, it consumes more area and power in comparison with conventional routers in two other structures. In addition, the high number of input ports becomes a performance bottleneck in terms of increased router complexity and contention probability inside the router (i.e. there are more input ports competing for an output port) [3]. Nonetheless, as the number of routers is decreased by the clustering approach in CIT, not only it reduces the area and power dissipation of the network but also the TSV area footprint is considerably diminished on each layer. Besides, the distance between two nodes of the same cluster in CIT is only one router so that the data transmission between nodes of the same cluster can be very fast. That is, the latency in CIT for distant nodes is more than that in the mesh-based 3D structures, but for close nodes the latency in CIT is lesser.

3.2 CMIT (Cluster Mesh Inter-layer Topology)

The structure of CMIT, depicted in Fig. 1(b), is basically similar to that of the mesh topology, except that for every layer the number of vertical channels has been reduced by sharing a vertical bus among routers of each cluster. That is, even preserving the advantage of the mesh on each layer, CMIT diminishes the number of inter-layer interconnections to meet constraints on the number of TSVs. In CMIT, each node (IP-core/memory) is connected to a router having at most 6 ports: one to the node, one to the bus, and four for the neighbors. Fig. 1(b) exhibits CMIT with 36 nodes, in which every four routers are grouped into a cluster on each layer. Even though CMIT achieves better area and power efficiency than the typical 3D mesh structure due to reducing the number of vertical channels, since several routers are connected to a shared vertical bus, the performance can be degraded when the inter-layer traffic is augmented. The arbiter should be placed in the middle layer of the chip to keep wire lengths as uniform as possible. The number of control wires of each arbiter increases with the number of nodes attached to the vertical channel (bus). As a result, the presence of a centralized arbiter is one of the reason why the number of vertical channels in the chip should be kept low [3][6]. We believe that, the proposed topology can keep the number of vertical channels low with a negligible performance penalty.

4. EXPERIMENTAL RESULTS

In this section, we compare the presented topologies with the conventional structures in terms of latency, power consumption, and area cost. We configure a 64-node (4×4×4) 3D stacked architecture. In this configuration, out of 64 nodes, 16 nodes are assumed to be processors and other 48 nodes are memories, i.e. DRAMs. In addition, three different 3D on-chip network topologies are considered for experiment: 3D hybrid structure, CIT, and CMIT. Each input port of the routers has two VCs where packets of different message types (request and response) are assigned to corresponding VCs to avoid message dependency deadlock [9]. The routers adopt the DOR routing and utilize wormhole switching. For all routers, the data width (flit) was set to 32 bits, and the buffer depth of each VC is five flits. Each inter-layer channel has its arbiter module and bus controller based on the dTDMA bus structure [3]. The presented configuration uses one flit for messages related to read requests and write responses, and the size of read request messages typically depends on the network size and memory capacity of the configured system. All the cores and routers are assumed to operate at 1 GHz. To estimate the power consumption of networks, we have used Orion library functions [10] as well as the power and delay values of vertical links in [7].

4.1 Physical Analysis

The number of routers and vertical channels in a chip affects the area and implementation cost. Thus, to compute the network area for each topology, we estimate the area of routers, cluster routers, and vertical channels. The network platform of each topology with the aforementioned configuration is synthesized with Synopsys Design Compiler using the UMC 0.09μm technology, while the

backend is performed with the Cadence Encounter tool. Depending on the technology and manufacturing process, the pitches of TSVs can range from 1µm to 10µm square [3][4]. In this work, the pad size for TSVs is assumed to be 5µm square with pitch of around 8µ, the flit-width is set to 32 bits, and each vertical channel requires 3×14 control wires for arbitration in a four stacked layer [3]. After the TSV area is calculated with the given values, the TSV footprint has been reduced from 0.41 mm2 in 3D hybrid to 0.1 mm2 in CIT and CMIT, resulting in about 75% area saving for TSV footprint. 3D hybrid occupies larger network area compared with CIT and CMIT, because each router in 3D hybrid has a compact transceiver module to interface with the vertical channel (bus) [3] and each bus should have its own arbiter module. Likewise 3D hybrid, the transceiver and arbiter modules are only integrated in cluster routers of CIT and CMIT. On the other hand, the total network area used by CIT is significantly lower than the other architectures (37% and 42% less than that of CMIT and 3D hybrid respectively) since the network is formed only by cluster routers.

4.2 Performance Analysis

To assess the performance of the presented topologies, the uniform and non-uniform synthetic traffic patterns have been considered separately. Eight burst sizes, from 1 to 8, are stochastically chosen according to the data length of the request. In the non-uniform mode, 80% of the traffic is local requests and 20% of the traffic is uniformly distributed to the non-local memory modules.

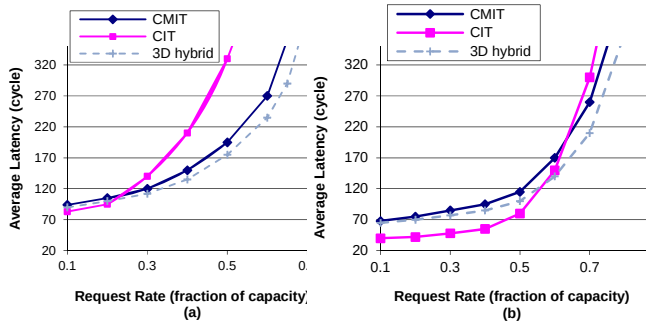

Fig. 2. Performance under (a) the uniform and (b) non-uniform traffic profiles.

Fig. 2(a) and (b) show the simulation results under the uniform and non-uniform traffic profiles, respectively. As demonstrated in Fig. 2(a), CIT has the lowest average latency in the low traffic load (<0.22). One of the foremost reasons for such an improvement is that CIT reduces the average hop count and improves load balance across the channels. But in high traffic load the performance of CIT is degraded considerably since the network bandwidth in CIT is lower than that of mesh-based structures. That is, the number of links in CIT is lower than that of mesh-based structures. Therefore, in the high traffic load, the traffic in CIT links is much higher than mesh-based structures. Also, the latency in CIT for distant destinations is significantly higher than that of mesh-based structures due to the high router complexity and contention probability, while for close destinations the latency of CIT is lesser. Thus, CIT might have better performance in such applications that most of requests are issued among neighboring nodes. This can be seen from the results in Fig. 2(b) where each processor sends requests to the memories based on the non-uniform traffic profile. CIT outperforms the others when the request rate is below the saturation point and most of the traffic is local. As a result, compared with the 3D hybrid and CMIT, the average latency of CIT is reduced by 20% and 30%, respectively.

4.3 Power Analysis

Using the simulator, the average power consumption of the presented topologies were calculated and compared under uniform and non-uniform traffic patterns close to the saturation point. According to Fig. 2 (a) and (b), the saturation points that have been considered for computing the average power values are 0.3 for uniform traffic and 0.5 for non-uniform traffic. The average power dissipation of the CIT network is 16% and 10% less than those of the 3D hybrid and CMIT schemes under the uniform traffic profile, respectively. The results indicate that the average power of CIT is 35% less than that of the 3D hybrid scheme and 30% less than that of the CMIT scheme under the non-uniform traffic profile. We can notice that although the power consumption of every cluster router is about 1.5 times higher than the power consumption of a typical router, the average power in the CIT network, compared to other schemes, is considerably lowered under non-uniform traffic profile since the average number of hops between two arbitrary nodes is less than the other presented schemes.

5. CONCLUSION

Since the area overhead of TSVs imposes constraints on the number of TSVs for existent 3D architectures, in this paper, two cluster-based topologies have been presented to deal with constraints on the number of TSVs. Experimental results have revealed that CIT and CMIT reduce the number of TSVs significantly and diminish the power consumption with low performance penalty under uniform traffic, but under non-uniform traffic, CIT reduces the average network latency too.

6. ACKNOWLEDGMENTS

The authors wish to acknowledge the Academy of Finland and Nokia Foundation for the partial financial support during the course of this research.

7. REFERENCES

[1] B. S. Feero, P. P. Pande, “Networks-on-Chip in a Three-Dimensional Environment: A Performance Evaluation,” IEEE Transactions on Computers, vol. 58, no. 1, pp. 32-45, Jan. 2009.

[2] M. Ebrahimi, et al., "Exploring Partitioning Methods for 3D Networks-on-Chip Utilizing Adaptive Routing Model," in Proc. of NOCS, USA, 2011.

[3] F. Li, et al., “Design and Management of 3D Chip Multiprocessors Using Network-in-Memory,” In 33rd International Symposium on Computer Architecture (ISCA), pp. 130–141, 2006.

[4] I. Loi and L. Benini, “An Efficient Distributed Memory Interface for Many-Core Platform with 3D Stacked DRAM,” in Proc. of the DATE Conference, Germany, pp. 99-104, 2010.

[5] I. Loi, F. Angiolini, L. Benini, “Supporting vertical links for 3D networks on chip: toward an automated design and analysis flow,” in Proc. Nanonets, 2007.

[6] S. Pasricha, “Exploring Serial Vertical Interconnects for 3D ICs,” in Proc. IEEE/ACM DAC, pp. 581-586, 2009.

[7] V.F. Pavlidis et al., “3D Topologies for Networks-on-Chip,” IEEE Transactions on Very Large Scale Integration Systems, 15(10):1081, 2007.

[8] Y. Qian et al., “From 2D to 3D NoCs: A Case Study on Worst-Case Communication Performance,” In. Proc. of the International Conference on Computer-Aided Design (ICCAD), pp. 555-562, 2009.

[9] S. Murali, and et al. “Designing message-dependent deadlock free networks on chips for application-specific systems on chips,” In Proc. VLSI-SoC, pages 158-163, 2006.

[10] A. Kahng et al., "Orion 2.0: A fast and accurate noc power and area model for early-stage design space exploration," in proc. of DATE, pp. 423-428, 2009.