Small-World Network Enabled Energy Efficient and Robust 3D NoC Architectures

Sourav Das, Dongjin Lee, Dae Hyun Kim, Partha Pratim Pande
School of Electrical Engineering and Computer Science
Washington State University, Pullman, USA
{sdas, dlee2, daehyun, pande}@eecs.wsu.edu

ABSTRACT
Three dimensional (3D) Network-on-Chip (NoC) architectures enable design of low power and high performance communication fabrics for multicore chips. In spite of achievable performance benefits, 3D NoCs are still bottlenecked by the planar interconnects. To exploit the benefits introduced by the vertical dimension, it is imperative to explore novel 3D NoC architectures. In this paper, we propose design of a small-world (SW) network based 3D NoCs. We demonstrate that the proposed 3D SW NoC outperforms its conventional 3D mesh-based counterparts. On average, it provides ~25% reduction in the energy delay product (EDP) compared to 3D MESH without introducing any additional link overhead in presence of conventional SPLASH-2 and PARSEC benchmarks. The proposed 3D SW NoC is more robust in presence of TSV failures and performs better than fault-free 3D MESH even in the presence of 25% TSVs failure.

Categories and Subject Descriptors
C.2.1 [**Computer-Communication Networks**]: Network Architecture and Design

General Terms
Algorithms, Performance, Design

Keywords
3D NoC; small-world; latency; energy dissipation; TSV

1. INTRODUCTION
Three dimensional (3D) Network-on-Chip (NoC) is an emerging paradigm that takes advantages of amalgamation of two emerging technologies, NoC and 3D IC. It allows for the creation of new structures that enable significant performance enhancements over more traditional solutions. With freedom in the third dimension, architectures that were impossible or prohibitive due to wiring constraints in planar ICs are now possible, and many 3D implementations can outperform their 2D counterparts. Existing 3D NoC architectures predominantly follow straightforward extension of regular 2D NoCs. However, this does not exploit the advantages provided by the 3D integration technology appropriately. The additional degree of freedom provided by the vertical connections enables design of more efficient irregular architectures. In this context, design of small-world network-based NoC architectures [1] is a notable example. It is already shown that either by inserting long-range shortcuts in a regular mesh architecture to induce small-world effects or by adopting power-law based small-world connectivity it is possible to achieve significant performance gain and lower energy dissipation compared to traditional multi-hop mesh networks [1][2]. In this work, we advocate that this concept of small-worldness should be adopted in 3D NoCs too. More specifically, the vertical links in 3D NoCs should enable design of long-range shortcuts necessary for a small-world network. These vertical connections give rise to close proximity among communicating nodes that would have been far apart in a solely planar system.

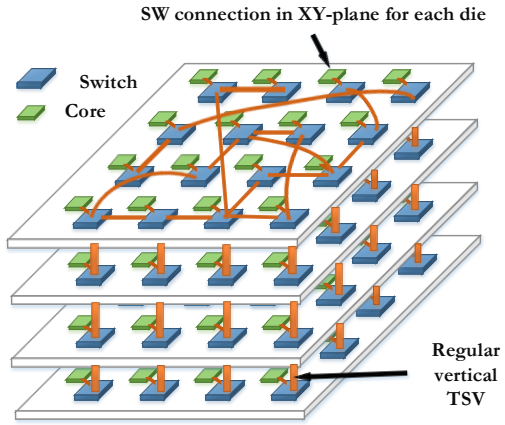

Figure 1: Conceptual view of 3D SW NoC with TSVs. For simplicity, only one logical XY-plane SW connection is shown here.

In this work, we propose design of 3D small-world (SW) NoC architectures. Considering the perfectly aligned vertical link placement constraints of Through Silicon Vias (TSVs), we first determine the suitable design parameters of 3D SW NoCs. Next, by exploiting the vertical dimension in a 3D IC, the tasks are mapped among the cores in such a way that physically long distant and highly communicating cores are placed along the vertical dimension, and hence overall system performance can be significantly improved. Through rigorous experiments we demonstrate that the proposed architecture is capable of achieving better performance and lower energy dissipation compared to conventional 3D MESH architectures. As the TSVs are costly in terms of fabrication, area overhead, energy dissipation; we investigate the performance degradation of different NoCs due to TSV failure. We demonstrate that for up to 25% of TSV failure, the proposed 3D SW NoC is capable of achieving better performance than a fault-free 3D MESH architecture.

2. RELATED WORK
3D NoC has emerged as one of the compelling solutions to design high performance and low power communication infrastructure for multicore chips. The natural extension of 2D planar architecture was the simple and regular 3D mesh-based NoC, which has been investigated in many existing works [3][4][5]. NoC-bus hybrid architecture was proposed in [6] that used central bus arbiter and

GLSVLSI '15, May 20 - 22, 2015, Pittsburgh, PA, USA.

http://dx.doi.org/10.1145/2742060.2742085

Dynamic Time Division Multiple Access (dTDMA) technique for bus access in the vertical dimension to reduce the network latency. To take advantage of vertical short distance inherent in 3D integration, 3D Dimensionally Decomposed (DimDe) NoC router architecture [7] was developed that reduces the total energy consumption, but latency was not minimized. Reducing the number of input ports, an improved version of 3D NoC router architecture was developed in [8]. As all of these architectures have buses in the Z-dimension, with increase in the network size, they are subject to traffic congestion and high latency under high traffic injection loads. Integration of 3D IC and on-chip photonics was exploited in the design of hybrid NoC architectures [9][10]. These architectures are capable of providing very high bandwidth and low power dissipation. However, the challenges of integrating two emerging paradigms, 3D IC and silicon Nano-photonics have not been overcome yet.

In this work we propose a robust 3D NoC architecture that combines the benefits of reduced vertical distance of 3D ICs and the low hop-count and robustness of small-world architecture. We present detailed design methodologies for developing the 3D SW NoC, evaluate its performance with respect to conventional 3D MESH and other irregular architectures by incorporating suitable network routing algorithms.

3. PROPOSED 3D NOC

Modern complex network theory provides a powerful method to analyze network topologies. Between a regular, locally interconnected mesh network and a completely random Erdös-Rényi topology, there are other classes of graphs, such as small-world and scale-free graphs. Small-world graphs have a very short average path length, defined as the number of hops between any pair of nodes. The average shortest path length of small-world graphs is bounded by a polynomial in log (*N*), where *N* is the number of nodes, making them particularly interesting for efficient communication with minimal resources [11]. NoCs incorporating small-world connectivity can perform significantly better than locally interconnected mesh-like networks [1], yet they require far fewer resources than a fully connected system.

3.1 Topology of Network

Our goal is to use the small-world (SW) approach to build a highly efficient 3D-NoC with planar links and TSVs in the vertical dimension. We consider 4 layers in this 3D NoC topology and the dimension of each layer is 10 mm x 10 mm.

3.1.1 Communication path length (μ)

In the proposed 3D NoC topology (Figure 1), each core is connected to a switch; the switches are interconnected using planar and TSV links. The topology of the NoC is a small-world network where the links between switches are established following a power law distribution. More precisely, the probability $P(i,j)$ of establishing a link between two switches i and j, separated by an Euclidean distance ℓ_{ij}, is proportional to the distance raised to a finite power as in [11]:

$$P(i,j) = \frac{\ell_{ij}^{-\alpha} f_{ij}}{\sum_{\forall i}\sum_{\forall j} \ell_{ij}^{-\alpha} f_{ij}} \quad (1)$$

The frequency of traffic interaction between cores, f_{ij}, is also factored in, so that the more frequently communicating cores have a higher probability of having a direct link between them. This frequency is expressed as the percentage of traffic generated from core i that is sent to core j. This approach implicitly optimizes the network architecture for a non-uniform traffic scenario. Getting now into details, the parameter α governs the nature of connectivity; in particular, a larger α would mean a locally connected network with a few, or even no long-range links. By the same token, a zero value of α would generate an ideal small-world network following the Watts-Strogatz model [11] – one with long-range shortcuts that are virtually independent of the distance between the cores. It has been shown that α being less than $D + 1$, D being the dimension of the network, ensures the small-worldness [12]. Overall, the parameter α, affects the NoC performance significantly. Thus, for our proposed architecture, we first focus on determining the parameter 'α' considering the constraints of the 3D network structure. To be specific, as we consider TSVs for vertical connection, these links need to be perfectly aligned along the z-dimension and we can apply irregularities following power law-based interconnection in the planar dies only.

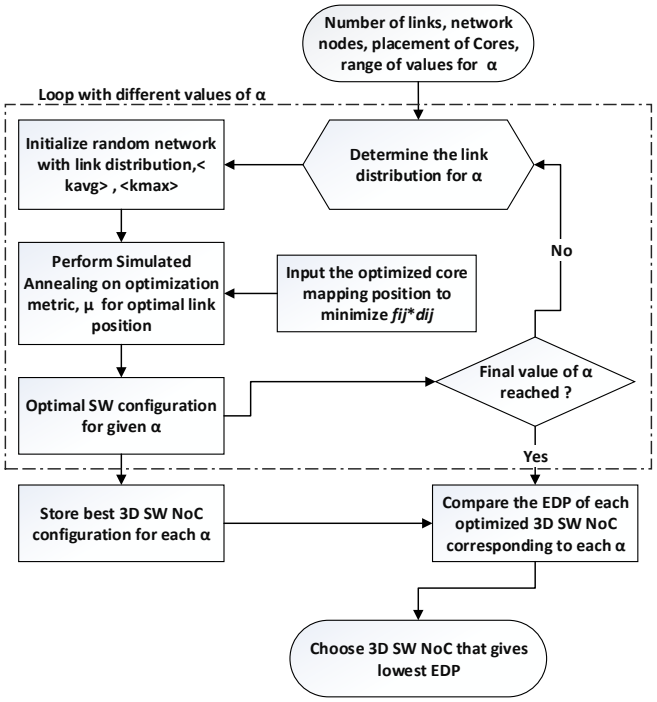

Figure 2: Our network creation algorithm for 3D SW NoC.

In order to determine α, we focus on optimizing the parameter, μ, which is the average path length for any message. Optimizing the average path length ensures less utilization of network resources and improvement in the network performance both in terms of latency and energy dissipation. We define μ as the product of hop count, frequency of communication and link length between any pair of source and destination-

$$\mu = \sum_{\forall i}\sum_{\forall j} (m * h_{ij} + d_{ij}) * f_{ij} \quad (2)$$

Where h_{ij} is hop count between switches i and j; m is number of stages inside a NoC switch; f_{ij} is frequency of communication and d_{ij} is physical distance corresponding to each hop calculated along the path. So any network that possesses low μ will achieve low message latency and energy dissipation and hence low energy-delay-product (EDP).

3.1.2 3D SW NoC design steps

As mentioned above, the parameter α governs the nature of the connectivity. Hence, we first determine the value of this parameter that will help us in designing the 3D SW NoC architecture with least EDP. To build 3D SW NoCs with optimum value of α, we develop the design flow shown in Figure 2. The inputs to the flow

are the total number of links, the total number of nodes, the locations of the cores, and α. The design flow is as follows:

1. For a given α, we calculate the link length distribution following (1). In this work, we constrain the number of total links equal to that of 3D MESH and the system size is 4x4x4.

2. We follow the constraint of perfectly aligned regular vertical link placement along the z-dimension. This is because along the z-dimension, we use TSVs and they need to be perfectly aligned.

3. The small-world network has an irregular connectivity. Hence, the number of links connected to each switch is not a constant. For fair comparison between our small-world network and 3D MESH, we assume that both of them use the same average number of connections, $<k_{avg}>$ per switch. This also ensures that the 3D SW NoC does not introduce additional links compared to 3D MESH. For a 64-core system, $<k_{avg}>$ is 4.5 considering all the switches including the peripheral ones. In addition, the maximum connectivity per node, $<k_{max}>$, is set to be 7 for the SW network as found in [13] [14].

4. To develop our SW network, we consider the communication frequency f_{ij} between any pair of source and destination nodes and try to optimize μ for network performance. We map the tasks among the cores in such a way that the overall μ decreases. To do this, we place highly communicating and long distant nodes along the vertical dimensions. Placing the physically long links in the vertical dimension provides two benefits. The wireline energy consumption reduces due to the reduction of the link length and the number of repeaters needed in a long planar wire. Moreover, the reduced network latency minimizing the probability of traffic congestion eventually reduces the overall power consumption.

5. After task remapping, we first build a random network with the link distribution determined at step 1. Then simulated annealing (SA) is performed to reduce μ. A solution perturbation method we use in the simulated annealing is to randomly choose an existing link, remove the link, and create a new link of the same length between two other nodes. The convergence criteria are the total number of perturbations, the lowest temperature, and the lowest bound (0.1% in our experiments) on the amount of performance improvement compared to the previous network configuration after a pre-determined number of perturbations. In this way, we develop the optimum 3D SW NoC configuration for a given α.

6. We build the optimum 3D NoC using the above steps for a given α. Consequently, we vary α within a certain range (determined experimentally in section 4) and determine the optimum 3D NoC for each α using our design flow. Finally to choose the optimum value of α for 3D SW NoC, we consider the particular network configuration that minimizes the EDP of the NoC. The simulation results will be shown in Section 4.

Following the above-mentioned steps and considering the constraints of network resources, we can develop optimized 3D SW NoC for any given set of applications.

3.1.3 3D SW- Bus NoC

In addition, we also consider 3D SW-BUS NoC architecture where we replace the network connections (Point-to-Point (P2P)) between the planar layers with TSV bus. As the distance between the individual 2D layers in 3D IC is small, bus is also a possibility for communicating in the Z-dimension [6]. For consistency with [15], our analysis considers the use of a dynamic time-division multiple-access (dTDMA) bus, although any other type of bus may be used as well.

3.2 Other 3D NoCs under Consideration

To benchmark the performance of our proposed 3D SW NoC architecture with respect to other irregular 3D NoCs, we consider the Mesh-Random-Random-Mesh (*mrrm*) and Random-random-random-random (*rrrr*) architecture as suggested in [15]. Like the 3D SW architecture, in this case also we keep the total number of links to be equal to that of a 3D MESH.

The characteristics of these two architectures are:

mrrm: Point to point (P2P) TSV based 3D NoC that consists of two planar layers of mesh-based interconnection architecture and the rest two layers have random interconnection patterns.

rrrr: P2P TSV based 3D NoC that consists of four layers of random interconnection networks.

To construct the *random* architectures for *mrrm* and *rrrr* configurations, 1000 random connection matrices were generated and we obtained the average of the average hop count over these 1000 matrices. Then the particular connection matrix, which has the closest average hop-count to the average of average hop-count, was selected for the *random* architecture

3.3 Routing Algorithm

3.3.1 ALASH

For regular 3D MESH architecture, XYZ or adaptive-Z are the preferred routing algorithms for their simplicity. For irregular architectures such as the small-world network, the topology agnostic Adaptive Layered Shortest Path Routing (ALASH) algorithm is proved to be suitable [16]. ALASH is built upon the layered shortest path (LASH) algorithm, but has better flexibility by allowing each message to adaptively switch paths, letting the message choose its own route at every intermediate switch. We incorporate the ALASH routing for our 3D SW architecture.

3.3.2 Elevator-First (EF)

Recently Elevator-first (EF) algorithm was proposed for 3D NoCs with limited vertical links [17]. To achieve deadlock-free routing in this algorithm, a message is restricted to revisit any xy-layer after it has already left that layer. To ensure this condition for an irregularly-connected xy-layer, we need to make sure that there should be at least one path between every source and destination pair in that layer. This puts additional constraints on irregular networks and SW NoC. Our proposed 3D-SW NoC has an irregular topology in the xy-plane, so we put this additional restriction on the proposed architecture keeping all other design parameters same. The performance of 3D SW NoC with EF algorithm is marked as 3D EF in this paper. We then compare the 3D EF's performance with respect to the ALASH-based design.

4. EXPERIMENTAL RESULTS AND ANALYSIS

4.1 Simulation Setup

To evaluate the performance of the proposed 3D SW NoC, we use a cycle accurate NoC simulator that can simulate any regular or irregular 3D architecture. Our system consists of 64 cores and 64 switches equally partitioned in four layers. Four SPLASH-2 benchmarks, FFT, RADIX, LU, and WATER [18], and five PARSEC benchmarks, DEDUP, VIPS, FLUIDANIMATE, CANNEAL, and BODYTRACK [19] are used. The width of each link is the same as the flit width, which is 32 bits. Each packet consists of 64 flits. The NoC simulator uses switches synthesized from an RTL level design using TSMC 65-nm CMOS process in Synopsys™ Design Vision. All switch ports have a buffer depth of two flits and each switch port has four virtual channels in case of

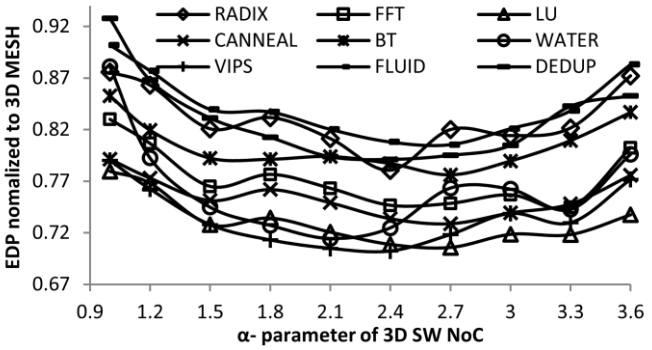

Figure 3: Variation of EDP w.r.t the connectivity parameter α.

irregular NoC. Hence, four layers are created in ALASH routing. Energy dissipation of the network switches, inclusive of the routing strategies, were obtained from the synthesized netlist by running Synopsys™ Prime Power, while the energy dissipated by wireline links was obtained through HSPICE simulations, taking into consideration the length of the wireline links.

4.2 Effects of Connectivity Parameter –α and Task remapping

In this section, we first determine the connectivity parameter- α for our proposed 3D SW NoC. We vary the connectivity parameter α from 1.0 to 3.6 and find the best α that gives us the lowest energy-delay product (EDP). Figure 3 shows the EDP values of our 3D SW NoC architecture for each α. The EDP values are normalized to those of 3D MESH NoC for all the benchmarks. As seen from the figure, as α increases from 1.0 to 2.4, the EDP decreases almost steadily. When α is small, large number of long-range links are likely to be placed at the cost of reduced local, short-range links. Although this configuration improves long-distance data exchange, it compromises the local communications, which have larger impact on the EDP than the long-distance communications. Thus, the EDP is high when α is small. As α increases from 2.4 to 3.6, most of the links become local, short-range links and the network approaches the regular multi-hop pattern that also results in high EDP. Therefore, we need to choose a value of α that achieves a compromise between these two extreme cases. If we increase α beyond 3.6, the long-range links become almost non-existent and the NoC architecture becomes very close to 3D MESH. Hence, we do not consider α beyond this.

To visualize the effect of α more clearly, Figure 4 plots the average EDP of the nine benchmarks. As shown in the figure, α=2.4 gives us the lowest EDP on average, so we choose α= 2.4 for our 3D SW NoC design.

Remapping tasks among the cores helps to improve NoC performance. However, task remapping may affect the traffic congestion. Figure 5 shows the percentage of total traffic carried by

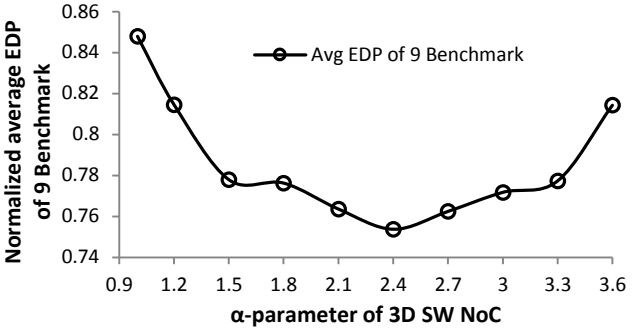

Figure 4: Average normalized EDP of the 9 benchmarks of 3D SW NoC with variation in α (normalized to 3D MESH).

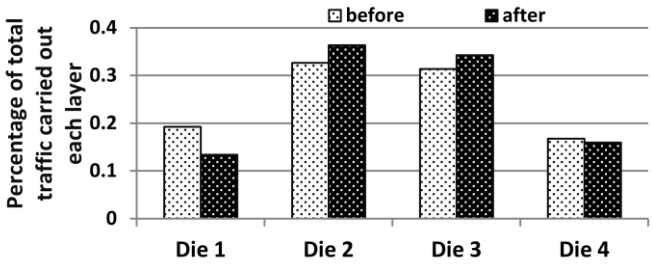

Figure 5: Percentage of total traffic carried out by each layer before and after task remapping.

each die before and after remapping the tasks. We observe from this figure that the task remapping slightly increases the traffic carried by the second and the third dies, but it reduces the traffic in the first and the fourth dies. Therefore, it is clear that our task remapping does not lead to noticeable traffic congestion in any particular layer.

4.3 Comparison with Other Existing NoCs

In this section, we compare the performance of the proposed 3D SW NoC with other 2D and 3D NoC architectures. Especially, we consider the following existing architectures: 2D MESH, 2D small-world (2D SW), and 3D MESH, 3D mesh-bus hybrid architecture (3D MESH-BUS), 3D small-world (3D SW), and 3D SW-bus hybrid architecture (3D SW-BUS). For regular connection of 2D and 3D MESH, we use the XY and XYZ- dimension order routing whereas for all the SW architectures, we use ALASH-based routing. We also implemented the Elevator-first routing algorithm for 3D-SW and 3D SW-BUS NoCs to demonstrate the benefits of ALASH routing algorithm for irregular architectures. The 3D SW and 3D SW-BUS with the Elevator-first algorithm are marked as 3D EF and 3D EF-BUS, respectively. Since we are interested in both network latency and energy consumption, we use the unified EDP for the NoC performance metric. Figure 6 shows energy-delay-product (EDP) values of all the NoC architectures. All the EDP values are normalized to those of the conventional 2D MESH.

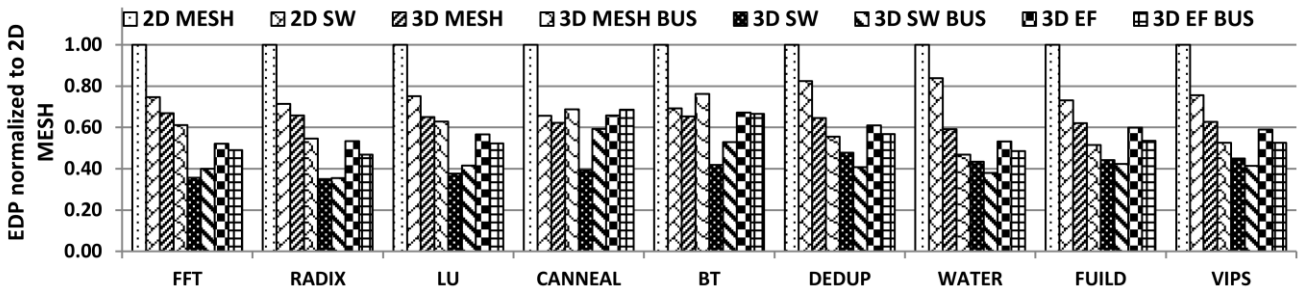

Figure 6: Energy-Delay product (EDP) normalized to 2D MESH for different NoCs.

The EDP for all the 3D NoCs is less than that of the 2D architectures. The value of the communication path length μ as shown in Table 1 helps us to analyze the reduction of EDP for 3D NoCs. Any network with lower value of μ, offers reduced weighted average hop count. On average, messages consume less network resources both from latency and energy perspective for that network. All the 3D NoCs have lower value of μ compared to 2D MESH and 2D SW and hence all of them perform better than 2D NoCs. 3D SW NoC possesses lowest value of μ among all NoCs. For 3D SW NoC, the remapping of the tasks for highly communicating and distant cores along the Z-dimension helps reducing the value of μ from 25.8 to 20.7. As a result, after remapping, the EDP value for 3D SW NoC reduces significantly. On an average, 3D SW shows ~25% reduction in EDP value compared to 3D MESH. Among the different 3D NoCs, 3D SW-BUS shows comparable EDP with 3D SW for relatively lower injection benchmarks i.e. LU, FFT, and VIPS. The bus is shared medium and performs efficiently whenever there are no access contentions among the connected nodes. For low traffic injection scenario, bus works as long range shot cuts between distant nodes without giving rise to contentions. However, for the comparatively higher traffic injection rates e.g. CANNEAL and BODYTRACK, bus access becomes a bottleneck and it results in traffic congestion. So network latency increases for 3D SW BUS and the performance degrades compared to 3D SW.

Next, we compare the performance of ALASH and Elevator-first routing algorithms implemented for the 3D SW and 3D SW-BUS architectures. ALASH routing algorithm ensures shortest path between a source-destination pair. Moreover, ALASH algorithm chooses paths based on traffic density and avoids traffic congestion on switches [14]. As a result it helps minimizing the latency and consequently energy consumption for any NoCs and performs better than other routing algorithms. On the other hand, Elevator-first routing algorithm is suited for 3D NoC with limited number of vertical links; it does not take into account traffic density of the X-Y planes. As a result, on average it encounters more traffic congestion and ultimately reroutes through longer paths compared to ALASH. Hence, by incorporating ALASH-based routing the latencies of 3D SW and 3D SW-BUS are less than 3D EF and 3D EF-BUS NoCs.

4.4 Comparison of 3D SW with Die Level Irregular NoCs

In this section, we compare the performance of the 3D SW NoCs with two recently proposed 3D irregular architectures, mesh-random-random-mesh (*mrrm)* and random-random-random-random (*rrrr)* as described in Section 3.2. Figure 7 shows the EDP for these architectures normalized with respect to the EDP of the 3D MESH. As the figure shows, the 3D SW architecture outperforms the *mrrm* and *rrrr* architectures. The *mrrm* and *rrrr* architectures have random distribution of interconnects instead of the power-law-based distribution. For *rrrr*, the link distribution contains large number of medium- and long-range links with the expense of short range links. As a result, the nearby communication suffers from long latency and power dissipation. In *mrrm*, the link distribution is very close to 3D SW NoC in two layers only whereas the rest of the two layers consist of multi-hop mesh links. The power-law-based 3D SW NoC architecture balances both nearby and long distant communication and hence its performance is superior to both- *rrrr and mrrm*. In addition, long-range links are generally traffic attractors, so they lead to higher traffic congestion for some nodes if they are not optimized properly. Finally, as shown in table 1, the average communication path length μ of 3D SW NoC is lowest and hence it performs better than all these irregular architectures.

Table 1: Path length (μ) for different NoCs

NoC type	Communication path length (μ)
2D MESH	37.98
2D SW	34.38
3D MESH	27.72
3D EF	27.2
3D *mrrm*	26.96
3D *rrrr*	27.76
3D SW (w/o task remapping)	25.84
3D SW (w/ task remapping)	20.71

4.5 Robustness of 3D SW Architecture

The 3D NoC architectures considered here use TSVs for vertical communication. TSVs are subject to failures due to voids, cracks, and misalignment [20]. Thus, we analyze the performance of 3D NoCs for a broad range of TSV failure scenario-5%, 10%, and 25% of the total vertical links. When some of the TSVs fail, we remove them, so the resultant NoC becomes partially vertically connected as suggested in [17]. To evaluate the robustness of our 3D SW NoC architecture and the ALASH routing algorithm separately, we perform two-step experiments. First, we demonstrate the robustness of the 3D SW NoC architecture. We consider the *mrrm* architecture for comparison as it was found to be the best among all the other 3D architectures considered in this work. Then we focus on analyzing the role of the ALASH and EF routing algorithms in presence of TSV failure. The 3D SW NoCs incorporating ALASH and EF algorithms are marked as 3D SW and 3D EF respectively.

Figure 8 shows the EDP of 3D SW NoC, *mrrm* and 3D EF incorporating different TSV failures. The EDP is normalized to a fault-free 3D MESH NoC. The TSV failure scenario has been chosen randomly from one thousand trials. It is seen that, for all range of TSV failures, the EDP of 3D SW NoC is always better than that of *mrrm*. Presence of different irregular connectivity and power-law based link distribution on each die for 3D SW NoC helps increase its robustness in case of TSV failure. In addition, we observe that, 3D SW NoC with 25% of the total vertical links failure has a lower EDP than completely fault free 3D MESH. Moreover, both *mrrm* and 3D EF perform worse than 3D MESH in presence of the same amount of vertical link failure.

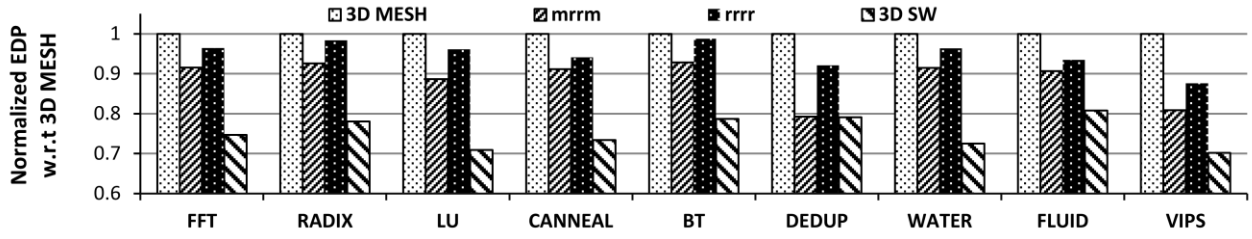

Figure 7: EDP value normalized to 3D MESH for different 3D NoCs.

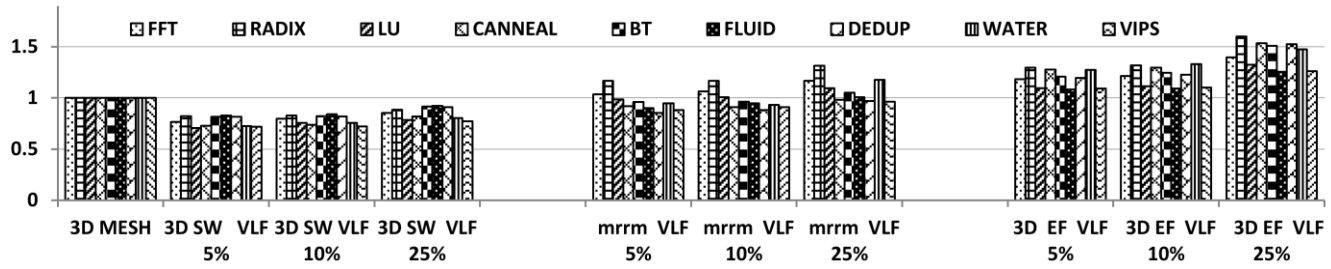

Figure 8: Normalized EDP in presence of random vertical link failure (VLF).

The Elevator-first algorithm mainly focuses on optimizing performance of partially vertically connected 3D NoCs [17]. Here we compare the performance of the NoC with vertical link failures by incorporating the ALASH and EF routing algorithms. ALASH ensures alternative shortest-path in case of link failure by adopting multi-path routing whereas EF always focuses on minimization of vertical distance without ensuring overall minimum path length. Consequently, removing higher number of TSVs from 3D SW NoC increases the traffic congestion probability for some of the vertical links in case of EF algorithm. As a result, the EDP of 3D EF with vertical link failure scenario is always higher than that of 3D SW implemented with ALASH.

5. CONCLUSION

This paper advocates design of small-world network-based 3D NoC architecture that predominantly uses the vertical links as the long-range shortcuts. In this way we are able to bring physically far and highly interacting nodes to logical proximity so that they are easily accessible to each other. This 3D NoC architecture not only outperforms existing regular mesh-based counterparts; it also offers better performance with respect to other recently proposed irregular architectures. The power-law based connectivity of this 3D SW NoC also makes it very robust against TSV failures. Even with 25% vertical link failure, it achieves better performance than fault-free 3D MESH-based NoC.

6. ACKONWLEDGEMENT

This work was supported in part by the US National Science Foundation (NSF) grants CCF-0845504, CNS-1059289, and CCF-1162202, and Army Research Office grant W911NF-12-1-0373.

7. REFERENCES

[1] U. Y. Ogras, R. Marculescu, ""It's a Small World After All": NoC Performance Optimization via Long-Range Link Insertion", IEEE Trans. on VLSI Systems, Vol. 14, pp 693-706, 2006.

[2] C. Teuscher, "Nature-inspired interconnects for self-assembled large-scale network-on-chip designs", e-print arXiv: arXiv: 0704.2852.

[3] V.F. Pavlidis, E. G. FriedMan, "3-D Topologies for Networks on-Chip," IEEE Trans on VLSI Systems, Vol 15, pp. 1081-1090, 2007.

[4] B. S. Feero, P.P. Pande "Networks-on-Chip in a Three-Dimensional Environment: A Performance Evaluation", IEEE Trans. on Computers, Vol 53, pp. 32-45, Aug. 2008.

[5] A.W. Topol, et al. "Three-Dimensional Integrated Circuits," IBM Journal of Research and Development, Vol. 50, nos. 4/5, July 2006.

[6] F. Li, C. Nicopoulos, T. Richardson, Y. Xie, V. Narayanan, M. Kandemir "High-Performance and Fault-Tolerant 3D NoC-Bus Hybrid Architecture Using ARB-NET Based Adaptive Monitoring Platform" IEEE Trans on Computers, Vol. 63, pp. 32-45, Aug. 2008.

[7] J. Kim et al., "A novel dimensionally-decomposed router for on-chip communication in 3D architectures," in Proc. of Annual Intl. Symp. On Computer Architecture, June 2007, pp. 138–149.

[8] A. M. Rahmani et al., "High-Performance and Fault-Tolerant 3D NoC-Bus Hybrid Architecture Using ARB-NET Based Adaptive Monitoring Platform", IEEE Trans. on Computers, vol 63, pp. 734-747, Mar 2014.

[9] D. Vantrease et al., "Corona: System Implications of Emerging Nanophotonic Technology," In Proc. of 35th Int'l Symposium Computer Architecture, pp. 153-164, June 2008.

[10] R. W. Morris, A. K. Kodi, A. Louri, R. D. Whaley"Three-Dimensional Stacked Nanophotonic Network-on-Chip Architecture with Minimal Reconfiguration", IEEE Transaction on Computers, Vol. 63, pp. 243-255, Jan. 2014.

[11] T. Petermann, P. D. L. Rios, "Spatial small-world networks: a wiring-cost perspective", e-print arXiv:cond-mat/0501420.

[12] T. Petermann, P. D. L. Rios, "Physical realizability of small-world networks", Phys. Rev. E 73, Feb 2006.

[13] R. Kim, G. Liu, P. Wettin, R. Marculescu, D. Marculescu, P. P. Pande, "Energy-Efficient VFI-Partitioned Multicore Design Using Wireless NoC Architectures", In Proc. of Intl Conf on, CASES 2014.

[14] P. Wettin et al., "Performance evaluation of wireless NoCs in presence of irregular network routing strategies," in Proc. of DATE, pp.1-6, 2014.

[15] H. Matsutani et al., "Low-Latency Wireless 3D NoCs via Randomized Shortcut Chips", in Proc. of DATE, 2014.

[16] O. Lysne, T. Skeie, S. A. Reinemo, I. Theiss, "Layered Routing in Irregular Networks", IEEE Transaction on Parallel and Distributed System, vol. 17, 2006, pp. 51- 65.

[17] F. Dubois, A. Sheibanyrad, F. Petrot, M. Bahmani, "Elevator-First: A Deadlock-Free Distributed Routing Algorithm for Vertically Partially Connected 3D-NoCs", IEEE Trans. on Computers, vol. 62, 2013, pp. 609-615.

[18] S. C. Woo, M. Ohara, E. Torrie, J. P. Singh, A. Gupta,"The SPLASH-2 programs: characterization and methodological considerations," In Proc. of International Symposium on Computer Architecture, 1995, pp. 24-36.

[19] C. Bienia, "Benchmarking modern multiprocessors," Ph.D. Dissertation, Princeton Univ., Princeton NJ, Jan. 2011.

[20] I. Loi, F. Angiolini, S. Mitra, L. Benini, "Characterization and Implementation of Fault-Tolerant Vertical Links for 3-D Networks-on-Chip", IEEE Trans. Computer Aided Design of Integrated Circuit and System, vol. 30, 2011, pp. 124-134.