Reliability and Performance Trade-offs for 3D NoC-Enabled Multicore Chips

Sourav Das, Janardhan Rao Doppa, Partha Pratim Pande
School of EECS, Washington State University
Pullman, WA, USA
Email: {sdas, jana, pande}@eecs.wsu.edu

Krishnendu Chakrabarty
Department of ECE, Duke University
Durham, NC, USA
Email: krish@ee.duke.edu

***Abstract*—Three-dimensional (3D) integration provides the benefits of better performance, lower power consumption, and increased bandwidth through the use of vertical interconnects and 3D stacking. The vertical interconnects enable the design of a high-bandwidth and energy-efficient small-world (SW) network-based 3D network-on-Chip (3D SWNoC) for massive multicore platforms. However, the anticipated performance gain of a 3D SWNoC-enabled multicore chip may be compromised due to the potential failures of through-silicon-vias (TSVs) that are predominantly employed as vertical interconnects. In particular, due to the non-homogeneous traffic patterns, heavily used TSVs may wear-out quickly and can also contribute to the wear-out of neighboring TSVs. As a result, the mean-time-to-failure (MTTF) of those TSVs will decrease, which will adversely affect the overall lifetime of the chip. In this paper, we address this traffic-dependent TSV wear-out problem in 3D SWNoC. We demonstrate that by employing an adaptive routing mechanism, we can improve the MTTF of 3D SWNoC significantly while still providing 21% lower energy-delay-product (EDP) compared to a conventional 3D MESH.**

I. Introduction

A small-world (SW) network-based 3D NoC (3D SWNoC) has been shown to outperform traditional 3D MESH and other irregular architectures in terms of latency and energy[1] [2]. In spite of this performance gain, all 3D architectures, including 3D SWNoC, suffer from potential performance degradation due to the failure of the through-silicon-vias (TSVs) that are employed as vertical interconnects [2]. Heavy workload is considered to be one of the main reasons for TSV failures [3][4]. Depending on the applications under consideration, workload experienced by the individual TSVs vary widely. Hence, their probabilities of failure, which directly depend on the workload, will also vary. In this work, our aim is to improve the overall reliability of the 3D SWNoC by incorporating an adaptive routing scheme to homogenize the traffic utilizations of all the vertical links (TSVs).

We make three key contributions in this paper. First, we study the effect of non-homogeneous traffic utilization on the mean-time-to-failure (MTTF) of different TSVs. Second, we analyze how the TSVs with smaller MTTF values affect the overall performance and hence, the lifetime. Finally, we develop an adaptive routing algorithm to maximize the lifetime of the critical TSVs while minimizing the performance penalty.

II. Related Work

Several 3D NoC architectures have been developed so far to improve the performance of multicore chips [5][6]. Among these architectures, small world network-based 3D SWNoC has been found to be highly energy-efficient [2]. Most of the 3D NoCs predominantly utilize TSVs as the vertical communication links. Prior work has analyzed the stress due to high current flow in TSVs and demonstrated that such stress can increase the resistance of TSVs due to electro-migration and may ultimately lead to TSV failure [3]. As a result, the performance of 3D NoC degrades over time. To address TSV failures, a fault-tolerant 3D NoC architecture is proposed in [6]. To increase the 3D IC lifetime and reliability, an in-field TSV repair framework was developed to improve the MTTF of TSV grids through judicious usage of spare TSVs [4]. However, all these methods focus on redistribution of the signals among the redundant TSVs and it is unclear how the workload of 3D NoC affects the individual TSVs. In addition, the fault tolerance and performance improvement of 3D NoC comes at the cost of extra TSVs, routing complexity, and significant area overhead. In this work, we study how the workload-induced stress on individual TSVs affects their MTTF values and the overall NoC lifetime. We also analyze how the EDP of the 3D NoC is adversely affected in such cases. Subsequently, we improve the reliability of 3D SWNoC by incorporating adaptive routing.

III. Problem Formulation

In this section, we explain the traffic-dependent TSV wear-out problem in a 3D SWNoC. Then, we describe our adaptive routing to increase the overall reliability of the 3D SWNoC.

A. TSV Wear-out Problem for NoC

The performance gain of 3D NoCs can be attributed to the presence of vertical links (generally TSVs) between the planar dies. As a result, the electrical characteristics of the TSVs have significant effects on the overall NoC performance. The RC characteristics of the TSVs change over time and increase non-linearly depending on the traffic load [3]. Under heavy traffic and high link utilization, the increase in resistance (R) of any TSV is more prominent than the change in capacitance [3]. The net effect of heavy TSV utilization is that it increases the delay of that particular TSV. In the worst case, the delay of the TSV can even increase more than one clock cycle period. This excessive delay is equivalent to '*TSV failure*'.

The '*wear-out*' effect of TSVs carrying heavy traffic load decreases the overall lifetime of the chip. The average time duration before a TSV fails is referred to as the mean-time-to-failure (MTTF). The effect of high TSV utilization can be modeled as increase in TSV resistance over time [3]:

$$R(t) - R_0 = A \ln\left(\frac{t}{t_0}\right), t > t_0 \quad \ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots\ldots (1)$$

where R_0 is the initial TSV resistance and mainly depends on

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.

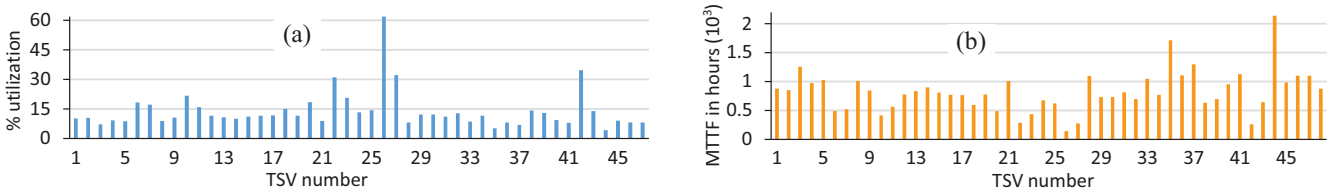

Fig. 1: Effect of non-homogeneous TSV utilization for 3D SWNoC with CANNEAL benchmark. (a) TSV utilization (in percentage of cycles that TSV actively transmit signals to total simulated cycles), and (b) MTTF values for individual TSVs

the fabrication process and component materials. Similarly, due to process variation, cracks and voids are created in the TSVs [3][4]. The parameter, t_0, is the time when the size of void becomes equal to or greater than the cross-sectional area of the TSV. The value of aging coefficient, A, depends on the amount of current that passes through TSVs in unit time. In this work, A, was taken to be 0.05 kΩ/log(s) based on the analysis reported in prior work [3]. The parameter $R(t)$ is the value of the resistance at time t in (1) and any increment in $R(t)$ will increase the TSV delay. When the resistance parameter, $R(t)$, increases to a value such that the delay of the TSV crosses the allowable limit, then that value of t is estimated to be the MTTF of that particular TSV. (Note that we estimate the MTTF in this manner because a more accurate computation of MTTF requires a statistically significant sample size of TSVs and actual measurement data.) The absolute value of this allowable delay depends on the performance requirement of the application under consideration. The parameter, $R(t)$, is also related to the traffic density and hence, the TSV utilization. The TSV utilization can be expressed as *utilized/simulated cycles* and by multiplying this ratio with the parameter t, we can calculate the traffic-induced TSV utilization.

We considered 2:1, 4:1 and 8:1 TSV serialization/ deserialization ratio and found that the 4:1 configuration provides the best trade-off between timing and area overhead. Hence, we considered a 4:1 TSV serialization/deserialization ratio in this work. Using the TSMC 28 nm libraries, the delay of the TSV was found to be 141 ps. This delay calculation includes the contribution from the serialization/deserialization circuit and the worst-case crosstalk noise from the neighboring TSVs. As mentioned above, $R(t)$ increases over time depending on the workload, which in turn increases the worst-case delay of the TSV. In this work, as an example, we set the allowable timing delay limit to be an additional 20% of the initial TSV delay. Once the TSV delay passes this limit, we assume that the TSV has failed. We consider the time to reach this 20% additional delay as the MTTF of that particular TSV.

B. *Addressing TSV Wear-out for NoCs via Adaptive Routing*

To increase the overall lifetime of the chip, first, we need to increase the MTTF of the TSVs those have very low MTTF values (i.e., they fail quickly). As mentioned above, in the 3D SWNoC, the main reason for lowered MTTF of any TSV is the heavy traffic and high TSV utilization. Hence, the traffic carried by the TSVs with very low MTTF value needs to be decreased. However, as the total amount of traffic for a particular application cannot be reduced, reducing the amount of traffic for a particular TSV will lead to an increase in the traffic experienced by other TSVs. As a result, the MTTF of those TSVs may decrease. For 3D NoC architectures with multiple planar layers, the traffic density in the middle layers is generally higher than that in the top and bottom ones. Hence, the TSVs between these layers experience much higher traffic density. As a result, the MTTF values of these TSVs are significantly lower, and ultimately, they influence the lifetime of the whole system. As an example, in Fig. 1, we show the traffic densities and the MTTF values of all the TSVs for a 64-core and four-layer 3D SWNoC for the CANNEAL benchmark (one of the PARSEC benchmarks with highest traffic injection load and skewed traffic). It is evident from Fig. 1(a) that the traffic densities of TSVs 17 to 32 are much higher than that of the others, and as expected, their MTTF values are significantly lower (Fig. 1 (b)). Our goal is to redistribute the traffic load in such a way that the lower MTTF values in the critical region will increase. We may have to sacrifice the TSVs with higher MTTF values to some extent. However, if the increase in MTTF for the critical TSVs is more significant than the reduction in MTTF of the non-critical TSVs then the average MTTF of the whole system will still improve. Redistribution of the traffic can be achieved by modifying the routing algorithm. Hence, we design an adaptive routing algorithm to enhance the MTTF of highly utilized TSVs.

1) *Adaptive Routing Algorithm*

The adaptive layered shortest-path routing (ALASH) algorithm is adopted for the 3D SWNoC as it represents the most suitable routing algorithm for this architecture [7]. ALASH is built upon the layered shortest path (LASH) algorithm [8], but it offers more flexibility by allowing each message to adaptively switch paths, and letting a message choose its own route at every intermediate router.

In order to avoid deadlocks, the network is divided into a set of virtual layers, which are created by dedicating the virtual channels from each router port into these layers. Then each source-destination pair is assigned to a layer such that the layer's channel dependency graph remains free from cycles. It is possible to induce deadlocks if a message switches back and forth between the layers [8]. Hence, a message is not allowed to revisit a layer to avoid deadlock.

To address the above-mentioned MTTF issue, we need to modify ALASH routing by updating the path taken by any message when the TSV utilization crosses a certain limit. To do that, every NoC router needs to have a counter to keep track of the traffic utilization of the TSV connected to it. For any region of the 3D SWNoC, the modified routing algorithm should aim to homogenize the TSV utilization pattern so that the MTTF of all TSVs in that region reaches a uniform higher level. If the utilization of any TSV crosses the average utilization of that region by one standard deviation amount (say a *pruning* rule), then the ALASH routing updates (say an *update* rule) the routing path by avoiding this TSV. We call this overall scheme as *prune-and-update,* where the pruning rule is used to prune the TSVs to be considered for routing and the update rule is

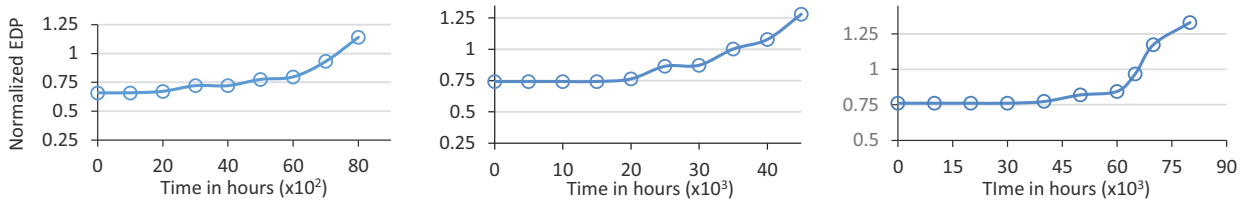

Fig. 2: Change in EDP of the 3D SWNoC over time due to TSV failures for different benchmarks of (a) high injection (CANNEAL), (b) medium injection (DEDUP), and (c) low injection (FLUID). EDP is normalized w.r.t. 3D MESH without any TSV failure.

used to select the alternate routing paths to minimize the performance penalty.

2) MTTF and Performance trade-off

To increase the MTTF value of highly utilized TSVs, the modified ALASH routing redistributes the traffic through the alternative paths. This redistribution may potentially increase performance penalty. The overall goal is to redistribute the workload in such a way that the performance degradation is minimal while the workload is distributed homogeneously among the TSVs in order to improve the overall MTTF value.

As mentioned above, the central region of any multi-layer 3D NoC is generally congested and TSVs present in that particular region carry a large proportion of the overall traffic. In this work, we consider a 3D NoC with four planar layers. Hence, TSVs connecting the second and third layers carry the highest amount of traffic. Redistributing the workload among the TSVs in this critical region may result in undesirable performance and energy overhead. Therefore, we need to determine the MTTF/performance/energy trade-offs. For a 3D SWNoC, the small-world architecture helps in reducing the performance degradation while improving the MTTF of highly utilized TSVs. The SW networks display remarkable resilience to high rates of link failures since the average distance between the nodes increase only by a small margin with the rate of failures. Hence, it compensates for the performance degradation when a TSV is avoided due to low MTTF value.

IV. Results and Analysis

A. Experimental Setup

To evaluate the performance of different NoCs, we use a cycle-accurate NoC simulator that can simulate any regular or irregular 3D architecture. Our system consists of 64 cores partitioned into four layers. The length of each packet is 64 flits and each flit consists of 32 bits. The routers are synthesized from an RTL level design using TSMC 28-nm CMOS process in Synopsys™ Design Vision. All router ports have a buffer depth of two flits and each router port has four virtual channels in the case of irregular NoC. The NoC simulator uses wormhole routing, where the data flits follow the header flits once the router establishes a path. For a regular 3D mesh-based NoC, XYZ-dimension order-based routing is used. For irregular architectures such as the SWNoC, the topology-agnostic ALASH algorithm is adopted [8]. The energy consumption of the network routers was obtained from the synthesized netlist by running Synopsys™ Prime Power, while the energy dissipated by wireline links was obtained through HSPICE simulations considering the length of the wireline links. We consider four SPLASH-2 benchmarks, namely, FFT, RADIX, LU, and WATER [9], and five PARSEC benchmarks, DEDUP, VIPS, FLUIDANIMATE (FLUID), CANNEAL, and BODYTRACK (BT) [10] in this performance evaluation.

B. Performance Degradation of 3D NoC due to TSV Failures

In Section III.B, we have shown the TSV utilization and the MTTF distribution for all TSVs for the CANNEAL benchmark. We have observed similar patterns for all the other benchmarks. In all cases, the TSVs placed in central region between the second and third layers of the four-layer 3D IC, experience high traffic load and hence, have lower MTTF value compared to the other TSVs. For the sake of brevity, we do not report the TSV utilization and MTTF values for the other benchmarks; instead we focus on examining the effect of workload-induced MTTF degradation of TSVs on the overall chip performance over time.

The occurrence of workload-dependent TSV failure will lead to performance penalty over time. Hence, we capture the EDP degradation as a function of time. In Fig. 2, we show the EDP vs time profile of the 3D SWNoC. As suitable representative cases, we choose the benchmarks with high (CANNEAL), medium (DEDUP), and low (FLUID) injection traffic. For a comparative study, the EDP is normalized to that of the 3D MESH NoC without any TSV failure. Note that the time on the horizontal axis considered in this case is the active data-transmit time, or in other words, the chip continuously runs for this period of time. We stopped collecting data when at least 50% of the total TSVs fail from the critical region and the performance of the 3D SWNoC becomes worse compared to the 3D MESH without any TSV failure.

From these figures, it is clear that the EDP degradation profile of 3D SWNoC for CANNEAL is much severe than the other two benchmarks, and its lifetime is order of magnitude

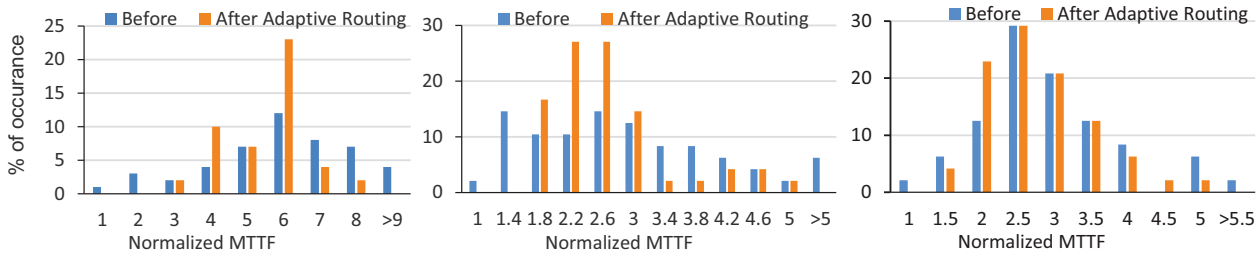

Fig. 3: Effect of adaptive routing on MTTF distribution for different traffic injection rate benchmarks. Distribution for (a) CANNEAL (high), (b) DEDUP (medium), and (c) FLUID (low). For each case, the MTTF is normalized w.r.t. lowest MTTF of the corresponding benchmark.

shorter than that for FLUID and DEDUP. This is expected as the TSV utilization factor is order of magnitude higher for CANNEAL compared to others and hence, the MTTFs of TSVs in this case are much lower than others. For DEDUP, the overall TSV utilization is almost double than that of FLUID. As a result, for FLUID, the 3D SWNoC maintains better EDP level for longer times compared to the other two benchmarks.

C. Adaptive Routing to Improve the MTTF value of TSVs

In this section, we analyze how adaptive routing increases the MTTF values of the TSVs in the critical region of the NoC. This will enhance the overall lifetime of the 3D NoC also.

1) Impact on TSV utilization and MTTF

To explain the effect of adaptive routing on TSV utilization, we consider three different benchmarks as before, viz., CANNEAL, FLUID, and DEDUP. With adaptive routing, the traffic utilization of critical TSVs decrease for all the benchmarks. As the total volume of traffic is constant before and after rerouting, the utilization of some low-utilization TSVs increases with adaptive routing. In other words, adaptive routing tries to homogenize the overall TSV utilization for both critical and non-critical regions.

Figs. 3 (a)-(c) show the MTTF distribution of all the TSVs before and after applying adaptive routing for CANNEAL, DEDUP, and FLUID respectively. To visualize the effect of adaptive routing, we have plotted the percentage of TSVs having a particular range of MTTF before and after adaptive routing. In all cases, the MTTF is normalized with respect to the lowest MTTF before applying adaptive routing.

From these figures, we observe that the MTTF distribution shifts towards the central region for all the benchmarks. Most importantly, the number of TSVs with the lowest MTTF value of the system before rerouting becomes zero and overall, the number of TSVs with low MTTF value decreases. Shifting the lower MTTF values towards the higher MTTF region implies that the adaptive routing enhances the lifetime of the most failure-prone TSVs. The factor of improvement is at least 2 (CANNEAL) and in some cases more than that. At the same time, the higher MTTF values decrease as well. However, these MTTF values correspond to the non-critical region and the overall lifetime of the chip is shorter than this high value. Therefore, shifting higher MTTF values towards the central region does not reduce the chip lifetime; rather, it increases it by increasing the MTTF value of critical TSVs through traffic redistribution. The adaptive routing algorithm redistributes the total traffic to increase the MTTFs for critical TSVs and thereby increases the overall chip lifetime.

D. Overall 3D NoC Performance with Adaptive Routing

In this section, we analyze the performance of the 3D SWNoC with adaptive routing algorithm. Fig. 4 plots the EDP of 3D SWNoC before and after incorporating the adaptive routing. For comparison, we plot the performance of 3D MESH as the reference and normalize all the EDP values w.r.t. it.

For all the benchmarks, incorporating the adaptive routing increases the EDP of 3D SWNoC from 7% to 13%. As expected, the 3D SWNoC with adaptive routing still shows an average of 21.8% lower EDP compared to the conventional 3D MESH. For adaptive routing, the improvement in the MTTF of individual TSVs, especially in the critical region is implemented through the redistribution of traffic. This approach, in some cases, utilizes longer path compared to the original shortest routing path of ALASH. As a result, network latency and energy consumption increases. In the worst possible scenario, this rerouting through longer paths can increase traffic congestion and overall high performance penalty. However, the proposed adaptive routing algorithm still uses ALASH, which is built upon the idea of finding the best shortest path and adaptive switching between virtual channels from the available resources. As a result, the average performance penalty for adaptive routing is only 8.5% while increasing the lifetime of the most critical TSVs by at least a factor of two and the overall lifetime of 3D IC significantly.

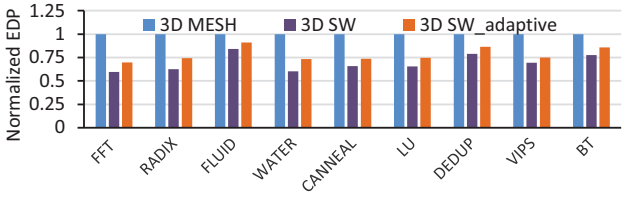

Fig. 4: Normalized EDP of 3D SWNoC (w.r.t. 3D MESH) with adaptive routing

V. Conclusion

We have analyzed and addressed the non-homogeneous workload-induced TSV wear-out problem in 3D NoC architectures. We have considered a small-world network inspired 3D NoC (3D SWNoC) as the test bed for this analysis as it significantly outperforms the conventional mesh-based architecture. However, the 3D SWNoC suffers from the lower MTTF value for certain TSVs due to the traffic patterns of the applications under consideration. By applying adaptive routing, we were able to improve the MTTF of the critical TSVs, which in turn improved the reliability of the overall system. The proposed adaptive routing scheme improved the overall MTTF of the 3D SWNoC significantly still maintaining 21% lower EDP compared to the conventional 3D MESH.

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] S. Das, J. R. Doppa, D. H. Kim, P. P. Pande, K. Chakrabarty, "Optimizing 3D NoC Design for Energy Efficiency: A Machine Learning Approach ", In Proc. of ICCAD, pp. 705-712, 2015.

[3] T. Frank et al. "Resistance increase due to electromigration induced depletion under TSV," In Proc. of Reliability Physics Symposium (IRPS), pp.3F.4.1-3F.4.6, 2011.

[4] L. Jiang; F. Ye; Q. Xu; K. Chakrabarty; B. Eklow, "On effective and efficient in-field TSV repair for stacked 3D ICs," In Proc. of DAC, pp.1-6, 2013.

[5] V. F. Pavlidis, E.G. Friedman, "3-D Topologies for Networks-on-Chip," IEEE Trans. on VLSI Systems, vol.15, no.10, pp.1081-1090, 2007.

[6] I. Loi et al., "A low-overhead fault tolerance scheme for TSV-based 3D network on chip links," In Proc. of ICCAD, pp.598-602, 2008.

[7] P. Wettin et al., "Design Space Exploration for Wireless NoCs Incorporating Irregular Network Routing," IEEE Trans. on CAD, vol.33, no.11, pp.1732-1745, 2014.

[8] O. Lysne, T. Skeie, S. A. Reinemo, I. Theiss, "Layered Routing in Irregular Networks", IEEE Trans. on Parallel and Dist. Sys, vol. 17, pp.51-65, 2006

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

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