

Inter-Tier Process-Variation-Aware Monolithic 3-D NoC Design Space Exploration

Shouvik Musavvir, *Student Member, IEEE*, Anwasha Chatterjee, *Student Member, IEEE*,
 Ryan Gary Kim, *Member, IEEE*, Dae Hyun Kim, *Member, IEEE*,
 and Partha Pratim Pande ^{id}, *Senior Member, IEEE*

Abstract—Monolithic 3-D (M3D) technology enables high density integration, performance, and energy efficiency by sequentially stacking tiers on top of each other. M3D-based network-on-chip (NoC) architectures can exploit these benefits by adopting tier partitioning for intra-router stages. However, conventional fabrication methods are infeasible for M3D-enabled designs due to temperature-related issues. This has necessitated lower temperature and temperature-resilient techniques for M3D fabrication, leading to inferior performance of transistors in the top tier and interconnects in the bottom tier. The resulting inter-tier process variation leads to the performance degradation of M3D-enabled NoCs. In this article, we demonstrate that without considering inter-tier process variation, an M3D-enabled NoC architecture overestimates the energy-delay-product (EDP) on average by 50.8% for a set of SPLASH-2 and PARSEC benchmarks. As a countermeasure, we adopt a process variation-aware design approach. The proposed design and optimization method distributes the intra-router stages and inter-router links among the tiers to mitigate the adverse effects of process variation. Experimental results show that the NoC architecture under consideration improves the EDP by 27.4% on average across all benchmarks compared to the process-oblivious design.

Index Terms—Energy-delay-product (EDP), monolithic 3-D (M3D), network-on-chip (NoC), performance, process variation.

I. INTRODUCTION

THREE-DIMENSIONAL (3-D) integrated circuits (ICs) have been shown to enable the design of high-performance and energy-efficient systems [1]. In particular, the network-on-chip (NoC) can heavily benefit from 3-D integration as the communication backbone of manycore systems. By taking advantage of a third dimension, 3-D NoCs provide a scalable communication fabric with lower hop-count, lower energy, and higher performance compared to their 2-D counterparts [2].

Manuscript received June 10, 2019; revised August 30, 2019 and October 21, 2019; accepted November 12, 2019. Date of publication December 18, 2019; date of current version February 25, 2020. This work was supported in part by the U.S. National Science Foundation (NSF) under Grant CNS-1564014 and Grant CCF 1514269, and in part by the USA Army Research Office under Grant W911NF-17-1-0485. (*Corresponding author: Partha Pratim Pande.*)

S. Musavvir, A. Chatterjee, D. H. Kim, and P. P. Pande are with the School of Electrical Engineering and Computer Science, Washington State University, Pullman, WA 99163 USA (e-mail: shouvik.musavvir@wsu.edu; anwasha.chatterjee@wsu.edu; daehyun.kim@wsu.edu; pande@wsu.edu).

R. G. Kim is with the Department of Electrical and Computer Engineering, Colorado State University, Fort Collins, CO 80524 USA (e-mail: ryan.g.kim@colostate.edu).

Color versions of one or more of the figures in this article are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TVLSI.2019.2954770

Modern 3-D integration processes have been widely adopted through-silicon-via (TSV) technology to connect planar dies together. However, there are several significant challenges to TSV-based 3-D integration. First, TSVs require additional fabrication steps like creating landing pads, wafer thinning, and bonding [3]. These fabrication and packaging-related challenges lead to lower yield rates and higher production costs for TSV-based designs [4]. Second, TSVs require a minimum keep-out-zone (KOZ) to reduce stress and coupling noise, introducing additional area overheads while undermining achievable integration density [5]. Third, even though TSV-based vertical links can improve communication in 3-D NoCs, they may fail due to crosstalk and electromigration [6].

Recently, monolithic 3-D (M3D) integration has been proposed as an alternative to TSV-based designs. In M3D designs, multiple tiers are processed sequentially on the same die [7] and monolithic inter-tier vias (MIVs) are used as vertical links instead of TSVs. The physical dimensions of MIVs ($\sim 50 \text{ nm} \times 100 \text{ nm}$) are several orders of magnitude smaller than TSVs ($1\text{--}3 \mu\text{m} \times 10\text{--}30 \mu\text{m}$) and are comparable to standard copper vias [8]. Similarly, the contact dimensions of M3D are much smaller ($\sim 100 \text{ nm}$ [9]) while TSV-based systems require contacts of $2\text{--}5 \mu\text{m}$. This allows us to achieve nanoscale contact pitch using M3D and attain the true benefits of vertical system integration. By facilitating nanoscale pitch, M3D enables us to examine gate- and block-level partitioning in circuits [7]. As a result, M3D offers much higher integration density and large reductions in total wire length over TSV-based counterparts. In addition, the direct wafer bonding technique in M3D achieves higher yields and lower costs compared to TSV-based integration [7], [10].

Naturally, NoC architectures can exploit the benefits of gate-/block-level partitioning in M3D integration by fabricating routers that span multitiers. In a recent study, M3D-enabled NoCs are shown to achieve 28% better energy efficiency compared to its TSV-based counterpart [11]. However, the investigations in [11] do not consider any M3D fabrication-related challenges or the benefit of lower interconnect capacitance from the reduced wire lengths in M3D designs [12].

While M3D architectures offer significant design flexibility and better energy efficiency compared to TSV-based designs, there are technology- and fabrication-related challenges that need to be addressed. M3D's sequential integration requires: 1) a low-temperature top-tier annealing process to prevent degradation in bottom-tier transistors [13] and 2) temperature-resilient tungsten interconnects in the bottom

tier to withstand high top-tier fabrication temperatures [7]. Unfortunately, these requirements degrade the transistors in the top tier and the interconnects in the bottom tier. Without considering these process-related effects it is likely that the performance and energy efficiency will be overestimated at design time.

In this article, to the best of our knowledge, for the first time, we demonstrate the importance of including these two M3D design requirements (tungsten interconnects at the bottom tier and a low-temperature top-tier annealing process) and the resulting inter-tier performance variation on the design and optimization of an M3D-enabled NoC architecture.

Our main contributions of this article are given as follows.

- 1) We demonstrate the necessity of including M3D process related effects in the design of NoCs. We show that a small-world NoC designed with M3D process parameters in mind lowers the energy-delay-product (EDP) with respect to the process-oblivious counterpart on average by 27.4% across all benchmarks under consideration.
- 2) We perform a detailed analysis of the effects of these M3D process-related parameters and the benefits of partitioning intra-router stages across tiers on the design of process-aware NoCs (i.e., intra-router stage placement and inter-router link distribution).
- 3) We find that the distribution of intra-router pipelined stages and inter-router links among the M3D tiers is strongly dependent on the values of the process variation parameters. We also find that the distribution of the intra-router stages and inter-router links depends on the benchmarks under consideration. All these observations show and justify the need for the M3D process-aware 3-D NoC design and optimization we propose in this article.

The rest of this article is organized as follows. The related work is discussed in Section II. Section III presents the challenges of M3D NoC design. In Section IV, we describe the problem setup and the proposed solution for the process variation induced performance variation. The experimental results are presented in Section V, and finally, the conclusions are provided in Section VI.

II. RELATED WORK

The merits of M3D-based designs have been explored in several works [14], [15]. M3D circuits provide reduced power and area compared to their 2-D counterparts. Motivated by the promise of monolithic integration, the CELONCEL design framework was proposed to explore the advantage of transistor/gate-level partitioning and cell-on-cell stacking design for M3D integration [15]. It was found that the footprint and wirelength of M3D-based designs are reduced by 37.5% and 16.2%, respectively, over their planar counterparts at the 45-nm technology node. This results in a 6.2% reduction of overall delay for M3D designs for the same technology node. Moreover, performance improves for more advanced technology nodes such as 22 [16], 14 [8], and 7 nm [17]. The effects of the number of planar tiers, tier-level partitioning, and MIV insertion methodology on the performance of M3D-based ICs were analyzed [18]. As the number of MIVs

in the design increases, the power saving improves as well. Researchers have also explored circuit design using transistor-level M3D integration [19]. While designing cores, they were able to reduce the area footprint and power consumption by 38% and 14%, respectively compared to the planar counterpart. Gopireddy and Torrellas [20] designed processors using M3D technology. They reported 25% improvement in performance and 39% less energy consumption compared to the 2-D counterpart. M3D systems enabled by nanotechnologies (N3XT) are proposed in [21]. N3XT employs recent nanotechnologies such as carbon nanotubes and M3D integration and achieves high-performance and energy efficiency. The reliability challenges of M3D-based designs have also been explored [22], [23]. This article in [22] demonstrated that M3D architectures can reduce circuit variance by tier partitioning. M3D circuits can endure nonideal effects like low frequency noise, bias temperature instability, and hot carrier injection and still meet the expected lifetime requirement [23]. Koneru *et al.* [24] demonstrated that the effects of electrostatic coupling and wafer-bonding defects can degrade the performance of M3D circuits if the interlayer dielectric thickness is less than 100 nm. Hence, the electrostatic coupling can be avoided by making the ILD thickness more than 100 nm.

Several works address the application of M3D technology to different types of circuits, e.g., 3-D field-programmable gate array (FPGA) [25], 3-D dynamic random access memory (DRAM) [26], and 3-D static random access memory (SRAM) [27]. It is demonstrated that by using M3D technology, we can reduce the total area, path delay, and power consumption of the circuit. Similarly, the authors in [11] explored the design space of 3-D NoCs and demonstrated the efficacy of M3D integration. Researchers also investigated voltage-frequency-island-based power management and the resulting thermal effects of M3D NoCs [28]. However, all these studies considered ideal characteristics for the transistors and interconnects (i.e., uniform delay and power consumption) across different tiers in the M3D ICs and ignored the effects of inter-tier process variation during design and evaluation.

Panth *et al.* [29] examined the effects of a realistic M3D fabrication process on the performance of M3D-enabled ICs. They found that the energy consumption and delay can increase significantly compared to an ideal M3D process. This article also evaluated the performance of M3D ICs that incorporate low-temperature annealing for top-tier transistors and tungsten interconnects in the bottom tier. Although M3D is a promising emerging interconnect technology, there is very little work on exploiting this to design a 3-D NoC. The work in [11] incorporates router partitioning in the design of an M3D NoC but it does not consider the performance benefit of multitier router stages or inter-tier performance variation between tiers. Therefore, our aim is to create a design process that integrates inter-tier process variation and how the inclusion of inter-tier process variation (and lack thereof) impacts the design and optimization of an M3D NoC.

III. CHALLENGES OF M3D-ENABLED DESIGNS

Although M3D enables higher integration density and better performance than TSV-based designs, fabrication challenges

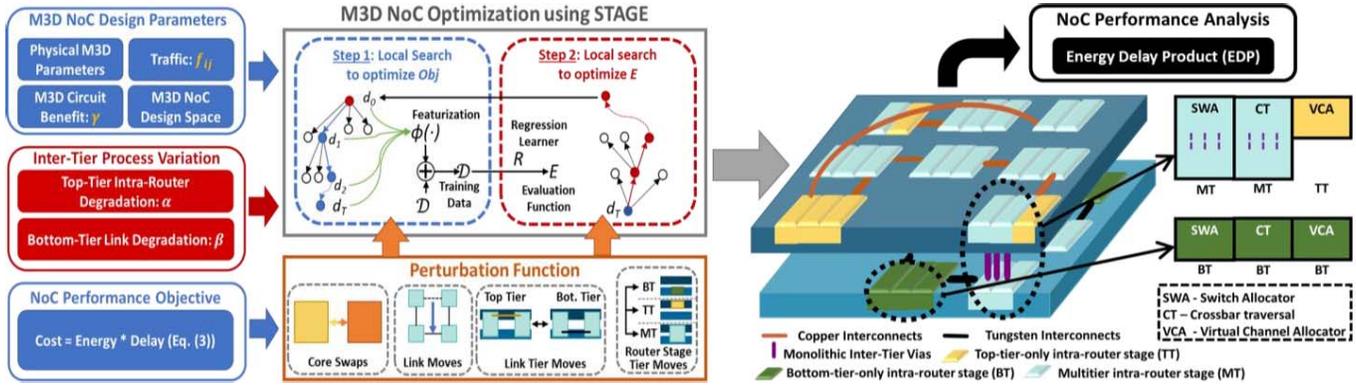


Fig. 1. Illustration of the process variation-aware design methodology. Implementation of the STAGE algorithm for M3D NoC optimization is shown in the second step. Step 3 shows an example optimized small-world network-enabled 9-node M3D NoC architecture. The legend indicates different components of the NoC.

need to be addressed. For example, fabricating upper tiers becomes a major challenge due to the sequential tier synthesis in M3D and the very thin inter-tier dielectric between the tiers. If ion implantation and annealing during top-tier fabrication use the standard thermal budget ($\approx 1050^\circ\text{C}$ [7]), the high temperature can damage the underlying interconnects and transistors in the bottom tier. As a result, two techniques have been proposed to realize top-tier transistors without damaging lower tiers: solid phase epitaxy regrowth (SPER) [30] and laser annealing [31].

In [7], it has been shown that temperature must be kept below 650°C to prevent any damage to the lower-tier transistors. This is accomplished by both SPER, the ion implantation step is done at temperatures as low as 600°C [30]; and laser annealing, upper-tier transistors can be fabricated while only heating up the bottom tier to 500°C [31]. However, both procedures have disadvantages. Transistors created using SPER have three times higher source-drain resistance compared to conventional transistors [30]. Similarly, transistors fabricated using laser-based annealing have 16%–28% lower on-current [13]. As a result, both processes introduce performance degradation for the top-tier transistors.

Unfortunately, although these temperatures are okay for lower-tier transistors, both SPER and laser annealing do not reduce the temperature enough to utilize copper back-end-of-line (BEOL) interconnects (the temperature must be kept within 400°C [7]). Therefore, a metal that can withstand higher temperatures such as tungsten is needed for BEOL interconnects in the bottom tier. However, tungsten has a higher resistivity than copper, which leads to inferior performance of bottom-tier interconnects [29].

These effects, degraded transistors in the upper tiers and higher resistivity interconnects in the lower tiers, can affect the design of NoCs. In particular, these inter-tier process variations cause nonuniform performance and energy consumption across the tiers. If these effects are not considered during design time, we may obtain overly optimistic latency and energy estimates and more importantly, suboptimal NoC configurations. This motivates our work into formulating a new M3D-variation-aware NoC design problem and optimization.

IV. PROBLEM FORMULATION AND OPTIMIZATION

In a manycore architecture, our goal is to place cores, routers, and links efficiently and design the intra-router stages and inter-router links optimally to achieve the best NoC performance. We begin by discussing the optimization goals for a standard 2-D NoC. Then, we extend this discussion to include M3D process variations and design considerations. Fig. 1 shows the overall flow of process variation-aware design of M3D NoCs. Finally, we present a framework that utilizes these optimization goals to design realistic M3D NoC architectures.

A. Network Latency and Energy Modeling of NoCs

In this section, we define general models for NoC latency and energy that can be applied to both 2-D and M3D systems. Later, we will provide specific details for 2-D (Section IV-B) and M3D (Section IV-C) systems. We model traffic-weighted NoC latency for a system with N cores and N routers as follows:

$$\text{Latency} = \sum_{i=1}^N \sum_{j=1}^N \sum_{r,u \in p(i,j)} \left(\sum_{m=1}^S \left(\left\lceil \frac{t_{r,m}}{15 \cdot \text{FO4}} \right\rceil \right) + \left\lceil \frac{l_u t_u}{15 \cdot \text{FO4}} \right\rceil \right) f_{ij}. \quad (1a)$$

For each iteration of the inner loop of (1a), determine x^*

$$x^* = \arg \min_{x \geq 0} \left\lceil \frac{\sum_{k=m}^{m+x} t_{r,m}}{15 \cdot \text{FO4}} \right\rceil (1 - \epsilon \cdot x). \quad (1b)$$

$p(i, j)$ gives the path between cores i and j (routers and links), where r and u are a router and a link along that path, respectively. The parameter $t_{r,m}$ is the intra-router stage delay for the m th router stage of router r with S router stages. It is important to note that $t_{r,m}$ depends on the number of ports and virtual channels (VCs) associated with the router. The parameters l_u and t_u are the length of the interconnect and unit length delay of the inter-router link u , respectively. $f_{i,j}$ is

the frequency of interaction between cores i and j . To capture the pipelining decisions made in a router, we assume a clock period of 15 fan-out-of-four (FO4) delays [2]. Therefore, the m th stage will take $\lceil (t_{r,m}/15 \cdot \text{FO4}) \rceil$ cycles. However, we can group up multiple stages to reduce the total cycles required by the router. Therefore, for each stage m , (1a) selects x^* consecutive stages [calculated by (1b)] following m that can still fit in $\lceil (t_{r,m}/15 \cdot \text{FO4}) \rceil$ cycles. The summation then considers stage $m + x^* + 1$ and repeats until all stages in the router have been considered. Here, ϵ is chosen to be small enough such that the arg min expression chooses the largest x where $\lceil (t_{r,m}/15 \cdot \text{FO4}) \rceil = \lceil (\sum_{k=m}^{m+x} t_{r,m}/15 \cdot \text{FO4}) \rceil$. Similarly, the link traversal latency is calculated as $\lceil (l_u t_u/15 \cdot \text{FO4}) \rceil$. Latency in (1a) captures the weighted sum of communication cost between a pair of cores.

Similarly, we model traffic-weighted NoC energy as follows:

$$\text{Energy} = \sum_{i=1}^N \sum_{j=1}^N \sum_{r,u \in p(i,j)} \left(\sum_{m=1}^S e_{r,m} + l_u e_u \right) f_{i,j}. \quad (2)$$

Here, for the network path between any two cores i and j , $e_{r,m}$ is the intra-router stage energy for m th router stage of router r with S router stages. The parameter e_u is the unit length energy of the inter-router link u .

In this article, our goal is to design energy-efficient and high-performance NoCs by minimizing traffic-weighted NoC latency and energy simultaneously. Therefore, we need a unified optimization metric for NoC latency and energy. To consider the effects of both latency and energy together, we use the product of traffic-weighted NoC latency and energy as the relevant performance metric for NoC optimization. Using (1) and (2), we represent the cost function for the NoC design optimization as follows:

$$\text{Cost} = \text{Energy} \cdot \text{Latency}. \quad (3)$$

Note that Cost is used only for the optimization. In the experimental results, we evaluate the final design using energy and latency values obtained from synthesis tools and full-system simulation. Further details are provided in Section V.

B. 2-D Router Delay and Energy Models

To determine the effects of M3D integration on the performance of an NoC, we first need to adopt an appropriate router model. In this article, we follow the virtual-channel router model proposed in [32]. However, it should be noted that any other router model can be adopted for this analysis ((1) and (2) are general in the number of stages). The virtual channel router has three pipelined stages, namely virtual channel allocator (VCA), switch allocator (SWA), and crossbar traversal [32].

The delay of each stage depends on the number of router ports (p), flit size (w), and the number of VCs (v). All these parameters in turn depend on the adopted NoC architecture. Their relationship is given in Table I [32]. For a regular NoC architecture like mesh, the delay of a particular pipelined stage will be the same in each router since each router has the same number of ports (except the routers at the edge). However, small-world networks have already been shown to

TABLE I
PARAMETERIZED DELAY EQUATIONS FOR
INTRA-ROUTER STAGES [32]

Intra-router stage (m)	Delay (FO4)
Virtual channel allocator (1)	$33 \log_4(pv) + 125/6$
Switch allocator (2)	$28 \log_4 p + 35/2$
Crossbar traversal (3)	$9 \log_8(w \lceil p/2 \rceil) + 6 \lceil \log_2 p \rceil + 6$

achieve significantly lower latency and energy consumption over mesh-based NoCs [11]. For these irregular small-world architectures, each router may have different number of ports. Hence, the delay of each pipelined stage varies depending on the router configuration. Using the model given in Table I, we determine the delay ($t_{r,m}$) of each pipelined stage (m) in every router (r) in terms of FO4 delay.

Energy consumption of the routers ($e_{r,m}$) depends on the capacitance of the logic cells and the interconnects of each stage [33]. For ease of notation, we will denote $t_{r,m}$ and $e_{r,m}$ as $t_{2D-r,m}$ and $e_{2D-r,m}$, for 2-D planar designs. Also, since 2-D systems use copper interconnects, $t_u = t_{Cu}$ and $e_u = e_{Cu}$ are the unit length delay and energy of transferring data through standard copper links, respectively. Using these models, we examine the effects of M3D integration on NoC design. It should be noted that both delay and energy models presented below are adopted exclusively for the optimization; we present the NoC performance evaluation methodology in Section V.

C. Process Variation-Aware M3D NoC Design

Although M3D integration allows us to build multitier routers, process variation from the M3D fabrication process causes the NoC to suffer from latency and energy penalties. As discussed in Section III, any router logic components in the top tier and inter-router links in the bottom tier will experience slowdowns due to process-related transistor degradation and the higher resistivity of tungsten, respectively. These effects could dominate the benefits obtained from tier partitioning, which would reduce the overall performance gain compared to its 2-D counterpart. Therefore, in addition to placing the links between routers, it is our objective to choose the tiers for each router stage and link of the M3D-based NoC to minimize the EDP. Hence, we consider the following design choices separately for each router stage and link (see Fig. 1). We restrict the tier partitioning of intra-router stages to two tiers following [34]. Although it is possible to stack more than two active tiers according to some monolithic 3-D manufacturing groups, no one has presented actual data for transistor and interconnect degradation with more than two tiers. Therefore, we do not have accurate data to perform the analysis for M3D NoCs with more than two tiers and focus only on two-tier monolithic 3-D integration in our article. Please note that this represents a high-level discussion; methods for determining delay and capacitance. will be discussed later in the results section (Section V-C).

1) *Bottom-Tier-Only Intra-Router Stage (BT)*: Since tungsten interconnects do not significantly affect the performance of logic gates [29] and there is no degradation

in the bottom-tier transistors, the performance of the logic gates in single tier (BT) is the same as in its 2-D counterpart. As the delay of the intra-router stages is dominated by the logic gates, the performance of BT does not change from the 2-D design. Therefore, $t_{r,m} = t_{2D-r,m}$ and $e_{r,m} = e_{2D-r,m}$.

2) *Top-Tier Only Intra-Router Stage (TT)*: As the transistors in the top tier are degraded, the FO4 delay of the router logic in TT will be larger than that in BT. To determine the delay of the router stage in TT, we need to determine the FO4 delay in TT in presence of the transistor degradation. Hence the parameters $t_{r,m}$ and $e_{r,m}$ will vary depending on the degradation of the transistor on-current (α). The intra-router delay for the router stage can be expressed as follows:

$$t_{r,m} = \frac{\text{FO4}_{\text{TT}}}{\text{FO4}_{2\text{D}}} \cdot t_{2\text{D}-r,m} \quad (4)$$

where FO4_{TT} is the degraded FO4 delay in the presence of α and $\text{FO4}_{2\text{D}}$ is the ideal FO4 delay for 2-D design (see Table I).

Transistor degradation will also increase the logic capacitance in the TT stage [35]. As energy is proportional to the capacitance in the router stage, the stage energy is expressed as follows:

$$e_{r,m} = \frac{C_{\text{TT},m}}{C_{2\text{D},m}} \cdot e_{2\text{D}-r,m} \quad (5)$$

where $C_{\text{TT},m}$ and $C_{2\text{D},m}$ are the total capacitance of TT and 2-D stage, respectively. $C_{\text{TT},m}$ comprises $C_{2\text{D},m}$ and the incremental logic capacitance due to α . It should be noted that the interconnect capacitance of the intra-router stage remains the same as the 2-D counterpart.

3) *Multitier Intra-Router Stage (MT)*: By using block- and gate-level M3D integration, the interconnect length can be reduced to a factor of $(1/\sqrt{T})$ for T -tier systems [12]. Therefore, although the change of logic capacitance is insignificant, the interconnect capacitance decreases by approximately $1 - (1/\sqrt{T})$ [12]. This results in an improvement of FO4 delay (denoted as γ) for a multitier design compared to the single-tier counterpart. In this article, we assume that the multitier stages are equally distributed among the bottom and top tiers. The M3D physical design work in [36] achieves an area skew of less than 10% while tier partitioning. Hence, our assumption of equal distribution of the multitier stage among two tiers is reasonable. Hence, the delay and energy for the router stages are

$$t_{r,m} = (1 - \gamma) \cdot \left(\frac{1}{2} \cdot t_{2\text{D}-r,m} + \frac{1}{2} \frac{\text{FO4}_{\text{TT}}}{\text{FO4}_{2\text{D}}} \cdot t_{2\text{D}-r,m} \right) \quad (6)$$

$$e_{r,m} = \frac{C_{\text{MT},m}}{C_{2\text{D},m}} \cdot e_{2\text{D}-r,m}. \quad (7)$$

Here, $C_{\text{MT},m}$ is the total capacitance of the multitier stage. The capacitance of the top tier logic of multitier stage will increase as mentioned in TT whereas the interconnect capacitance will decrease compared to the 2-D counterpart.

4) *Inter-Router Link Placement*: The delay and energy incurred by the inter-router links will depend on its tier placement. The inter-router links in the top tier use copper and do not suffer from any performance degradation (i.e., $t_u = t_{\text{Cu}}$ and $e_u = e_{\text{Cu}}$). On the other hand, the bottom-tier tungsten

interconnects exhibit higher resistance compared to the copper-based counterpart. Hence, the inter-router links in the bottom tier suffer from higher delay and energy consumption:

$$\begin{aligned} t_u &= t_W \\ e_u &= e_W \end{aligned} \quad (8)$$

where t_W and e_W are the delay and energy of tungsten links per unit length, respectively. Since tungsten has higher resistivity than copper, $t_W > t_{\text{Cu}}$ and $e_W > e_{\text{Cu}}$. We define the interconnect slowdown factor for tungsten interconnects as β , where β is the degradation in the propagation delay of the tungsten interconnects compared to its copper counterpart. As the resistivity of nanoscale tungsten wires changes based on the dimensions and geometry [29], β will vary too, which will be captured in (1) and (2) by t_W and e_W , respectively. Since these inter-router links are connected to the input and output stages of routers, a link and its respective router stages must be on the same tier. Therefore, during optimization (see Section IV-D), we constrain that top-tier links must be connected to TT or multitier router stages and bottom-tier links must be connected to BT or multitier router stages only.

D. M3D NoC Design Optimization

Since each router stage can be on different tiers (TT, BT, or multitier) with either top- or bottom-tier links, the design space complexity increases dramatically from a 2-D or TSV-based NoC to its M3D counterpart. For a system with N cores, M links, and S router stages, the upper bound of the design space increases to $3^{\text{SN}} 2^M N! \binom{(N(N-1)/2)}{M}$ where 3^{SN} corresponds to the tier selection for each stage in each router, 2^M corresponds to the tier selection for each link, $N!$ corresponds to core placement, and $\binom{(N(N-1)/2)}{M}$ corresponds to the link placement.

Such large search spaces make it difficult to utilize traditional optimization methods which rely on stochastic local exploration to reach the minima, e.g., simulated annealing (SA) or genetic algorithms (GAs). Therefore, intelligent search methods are necessary to reduce the run time and enhance scalability. We apply a machine-learning-based search technique, STAGE [37], to guide the search process. Prior works have already shown the efficacy of STAGE over SA and GA for different NoC architecture optimizations [11], [38].

STAGE works by utilizing past search trajectories to find better starting points. To accomplish this, STAGE iterates over two steps: 1) Hill climbing (or some other local search) to optimize Obj , the primary design objective and 2) Hill climbing to optimize E , a learned evaluation function that predicts how promising a design is as a starting point for Step 1. We show these steps in Fig. 1.

1) *STAGE Step 1*: Similar to simple hill climbing or SA, the first step attempts to minimize the target function Obj by making small steps (using a perturbation function S), accepting new designs if it reduces Obj , i.e., simple hill climbing. STAGE keeps track of all accepted designs in this search trajectory (d_0, d_1, \dots, d_T) and adds each design to a

data set D , as a pair $(\phi(d)Obj(d_T))$. Here, $\phi(\cdot)$ is a function that extracts pertinent features from the design.

2) *STAGE Step 2*: Using a regression learner, R , Step 2 learns the evaluation function $E(\phi(d)) = R(D)$. This evaluation function tries to predict the *Cost* of the final design of Step 1 starting from a particular design d . Ultimately, E can be used to predict the next best starting point for the search. Starting from the final design in Step 1, we use simple hill climbing to minimize E . This design is provided as the starting point to Step 1.

3) *STAGE Iteration*: STAGE iterates over Step 1 and Step 2 until the maximum number of iterations allowed ($Iter_{max}$) has been reached. After each iteration, we accumulate more data points in \mathcal{D} and learn a more accurate E which results in the search finding better designs. Our final output is the best design d^* with minimum Cost.

In the STAGE algorithm, the time complexity of the local search is $O(k.T)$, where k is the number of successors considered for each search node (i.e., local neighborhood) and T is the average number of greedy search steps before reaching the local optima from a given starting state. In each greedy search step, we evaluate each of the k successor nodes and select the best scoring successor node. It should be noted that the STAGE algorithm is run offline. Hence, energy consumption of STAGE is not included in the analysis and it has no area overhead.

In this article, we consider two different types of M3D NoCs, a process variation-oblivious M3D NoC that uses multitier for all router stages and uniformly distributes the links among layers [11] and our proposed process variation-aware M3D NoC. To utilize STAGE for M3D NoC design, we use the following definitions. For the perturbation function, S , we consider two types of perturbations to move to neighboring designs in all M3D architectures: 1) swapping the position of two cores and 2) moving a link between a pair of routers with another of the same length between two other routers. For process-aware M3D designs, we use two additional perturbations where: 1) a router stage is switched among the three different stage types (TT, BT, and multitier) and 2) a link is switched between the top and bottom tiers. For each call of the perturbation function, S , the function stochastically chooses one among the available perturbations described above. We show these perturbation choices in Fig. 1.

The feature selection ($\phi(\cdot)$) is an important aspect of STAGE as relevant features allow us to learn an accurate evaluation function, E . We use random forest regression to learn, E , however any other learner that is quick to evaluate and sufficiently expressive to fit the training data can be used to similar effect. Since the NoC performance directly depends on the average network hops and traffic-weighted hop-count, we use these metrics as features. Traffic-weighted hop-count is the sum of the products of hop count and communication frequency (f_{ij}) between routers i and j over all source-destination pairs. In addition, we use the clustering coefficient (C_c) for each router as a measure of a router's connectivity with its neighbors. The clustering coefficient captures the connectivity of one core with its neighbors. Although the hop count considers mainly long-range communication,

TABLE II
DIMENSIONS OF A TSV AND AN MIV [8]

Dimension	TSV	MIV
Diameter	4 μm	50 nm
Depth	20 μm	50 nm
Pitch	8 μm	44 nm

the clustering coefficient focuses more on local connectivity among the immediate neighbors. These three features, average network hops, traffic-weighted hop-count, and the clustering coefficient, are used as input to E . For process-aware M3D NoC designs, we use two additional features: bottom-tier inter-router and top-tier intra-router performance penalty to account for the process variation's effects on the M3D NoC design.

V. RESULTS AND ANALYSIS

In this section, we describe the experimental setup followed by the design considerations and architectural adjustments for our process-aware architecture. Then, we present a detailed performance analysis of the M3D NoCs with process variation. Finally, we provide allowable range for the process variation parameters that will make the process-aware design more energy-efficient than the TSV-based counterpart.

A. Experimental Setup

In this article, we consider a 64-core system where each core is associated with a dedicated router. We use GEM5 [39], a full system simulator, to obtain detailed processor- and network-level information. Using Gem5's full-system mode, we simulate $\times 86$ cores running Linux. We use the MOESI_CMP_directory cache coherence protocol. Each core consists of private 64-kB L1 instruction and data caches and a shared 8-MB L2 cache. We consider four SPLASH-2 benchmarks (FFT, RADIX, LU, and WATER) [40] and four PARSEC benchmarks [DEDUP, VIPS, CANNEAL (CAN), and FLUIDANIMATE (FLUID)] [41]. These benchmarks are selected because they vary widely in communication and computation patterns [42].

B. NoC Design and Baseline

For each router, we use the three-stage model shown in Table I [32]. Each router has four VCs ($v = 4$) per port. Each packet contains six flits and each flit consists of 32 bits ($w = 32$). From the M3D NoC optimization (see Fig. 1), we obtain the link placements (which gives us the number of ports on each router p), the tier placements for all router stages and links, and the core placement. The routers are synthesized from a register-transfer level (RTL) design using a TSMC 28-nm CMOS process in Synopsys Design Vision. For the final NoC experimental analysis, we pipeline each router and inter-router stages assuming a standard 15-FO4 delay per clock cycle [2]. The EDP here is the product of latency and energy of this pipelined NoC. Table II shows the physical dimensions of a TSV and an MIV. It should be noted that as the dimensions of MIVs are comparable to

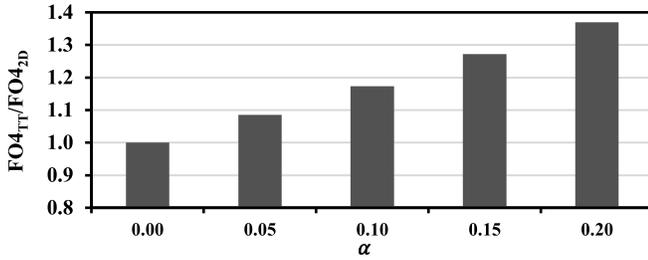


Fig. 2. FO4_{TT}/FO4_{2D} for different values of α .

that of local vias, using standard 2-D cells in the synthesis of the NoC routers does not add any additional overhead [8]. Energy consumption of the NoC links is determined using HSPICE simulations, taking their lengths and resistivity (the bottom tier uses tungsten, whereas the top tier uses copper) into consideration. We finally obtain latency and energy consumption values from full-system NoC simulations using SPLASH-2 and PARSEC benchmarks, the synthesized netlist, and HSPICE simulations.

For the NoC topology, we consider two cases: small-world network-enabled 3-D NoC (SWNoC) and traditional 3-D mesh. We create the SWNoC architecture by following a power-law-based link distribution as elaborated in [38]. It should be noted that we have already demonstrated in our prior works that SWNoC outperforms any other regular and application-specific 3-D NoC architectures [11], [38]. For 3-D mesh, we use standard XYZ-dimension-order routing. Since small-world networks are irregular topologies, we adopt the topology-agnostic adaptive layered shortest path routing (ALASH) for SWNoCs [43].

C. M3D Transistor and Interconnect Characteristics

In this article, we consider a two-tier M3D process. In order to deliver a thorough analysis of M3D NoCs, we must first provide a practical range for the degradation of transistor on-current (α), slowdown factor due to tungsten interconnects (β), and M3D improvement in FO4 delay (γ) parameters (see Section IV-C) and determine their effects on the design and optimization of M3D NoCs.

For α , prior work has reported a maximum of 20% degradation in the top-tier transistors [34]. Hence, we can consider $\alpha = [0.05, 0.20]$ in 0.05 increments. The parameter α affects FO4_{TT}, $C_{TT,m}$, and $C_{MT,m}$, and hence, the energy and delay of TT and MTs. We used Cadence Virtuoso to determine FO4_{TT} in the presence of α . Fig. 2 shows the ratio of FO4_{TT} to FO4_{2D} for different values of α . Here, $\alpha = 0$ is the case for nominal 2-D transistors. For every 5% increment in α , the FO4 delay degrades by approximately 9%. The increment of top-tier logic capacitance for TT [$C_{TT,m}$ in (5)] and multi-tier [$C_{MT,m}$ in (7)] is estimated from the relationship between transistor on-current and capacitance presented in [35]. In Fig. 3, the normalized capacitance for transistors is plotted for different values of α . Here, the capacitance is normalized with respect to a nominal 2-D transistor ($\alpha = 0$). As expected, the capacitance increases almost linearly with [35] α .

For β , the impact of tungsten interconnects in the bottom tier depends on the metal layer and technology node [29].

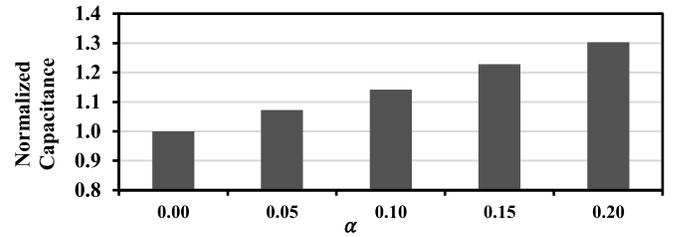


Fig. 3. Normalized transistor capacitance for different values of α . Capacitance is normalized with respect to the nominal transistor.

TABLE III
RANGE OF α , β , AND γ PARAMETERS

Parameter	Description	Range
α	Top-tier transistor degradation	0.05-0.20
β	Bottom-tier interconnect degradation	0.10-0.30
γ	Tier partitioning advantage in M3D	0.10-0.20

We consider tungsten interconnects in metal layers three through ten for a TSMC 28-nm process. For each layer, we characterize the delay of the tungsten interconnect by changing the resistivity [29] in Cadence Virtuoso. Our experimental analysis shows that the tungsten interconnects introduce 10%–30% additional propagation delay per unit length. So, we use $\beta = [0.10, 0.30]$ in 0.10 increments.

For γ , M3D designs achieve up to 25% improvement in clock frequency compared to their 2-D counterpart [44]. Thus, we use two values for γ (0.10 and 0.20) in this article. To determine the multitier stage energy in (7), we find $C_{MT,m}$, which must consider the increased logic capacitance at the top tier due to α [35] and the lower interconnect capacitance due to tier partitioning. We show the range of α , β , and γ values in Table III.

D. Router Stage and Link Distribution With Process Variation

Under ideal conditions ($\alpha = 0$, $\beta = 0$), naturally, all intra-router stages in an M3D NoC will be multitier and the links will be placed evenly in both tiers. Unfortunately, as discussed in Section IV-C, top-tier transistor degradation ($\alpha > 0$) and bottom-tier interconnect degradation ($\beta > 0$) cause significant slowdowns in the M3D NoC. Moreover, these slowdowns are not uniform across different tiers. In addition to these nonuniform effects of M3D process variation, the delay and capacitance of each intra-router stage vary with the number of bits per flit, ports, and VCs in the router (see Table I). Thus, to minimize the overall EDP for a given system configuration (number of bits per flit, number of ports, number of VCs, packet size, α , β , and γ), we should optimize the distribution of the intra-router stages and inter-router links. We present the analysis for SWNoCs followed by mesh-based NoCs.

E. M3D SWNoC Architecture Optimization

In SWNoCs, the distribution of intra-router stages and links depends on the process variation effects and traffic distribution.

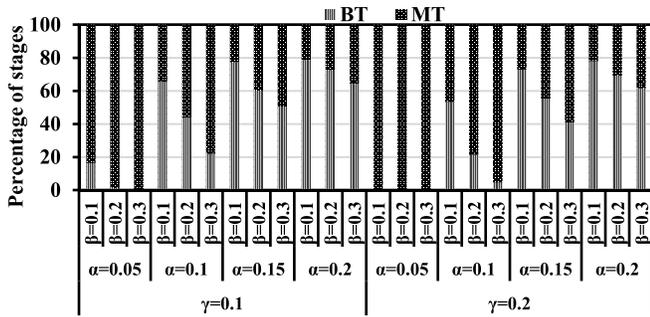


Fig. 4. Tier distribution of all intra-router stages for SWNoC (CAN).

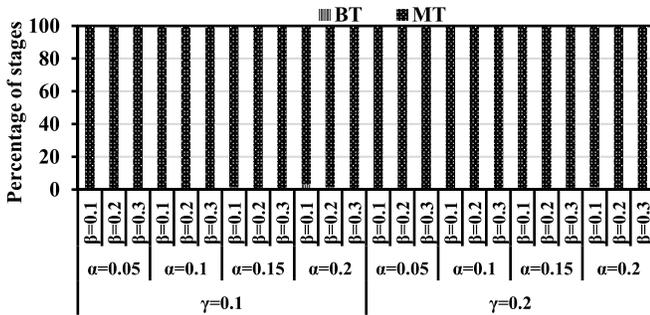


Fig. 5. Tier distribution of the crossbar stages for SWNoC (CAN).

Across a range of α , β , and γ , Fig. 4 shows the tier-wise distribution of all intra-router stages optimized for the CAN benchmark as an example. Since there is a trade-off between BT and multitier in terms of link versus logic degradation, we see different distributions of BT and multitier. Although the logic in the top tier of a multitier stage has longer delay and larger capacitance than the nominal values, the overall wirelength is shorter than sole BT stage. Therefore, some of the router stages can benefit from tier partitioning depending on α , β , and γ . For example, Fig. 4 shows that approximately 80% of the router stages are multitier and 20% are BT when $(\alpha, \beta, \text{ and } \gamma)$ is (0.05, 0.1, 0.1). However, if α increases, more stages are chosen to be BT to avoid the intra-router performance penalty originating from the top-tier transistor degradation as explained in (4) and (5). On the other hand, since γ improves the MT delay (6), there are more multitier stages at higher values of γ . Notably, the result shows that the SWNoCs do not have any TT stages. In multitier, only half of the logic cells suffer from transistor degradation, whereas all logic cells in TT experience transistor slowdown. Moreover, the speedup due to multitier logic (γ) results in the superior performance of multitier compared to TT. As a result, the optimization always chooses multitier over TT for all intra-router stages.

It should be noted that the distribution of different types of intra-router stages depends on their circuit composition. Since the crossbar stage is dominated by the interconnect capacitance [33] and interconnects are heavily reduced in multitier (interconnect capacitance decreases by 29.3% for the two-tier system), the energy saving in the interconnects offsets the transistor slowdown. In Fig. 5, it can be seen that besides a few crossbar stages at low values of β , all crossbar stages are multitier for every combination of α , β , and γ .

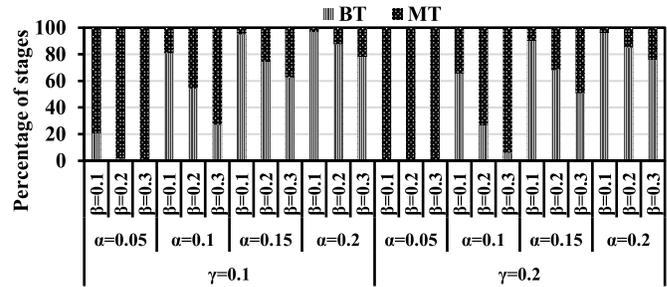


Fig. 6. Tier distribution of the SWA and VCA stages for SWNoC (CAN).

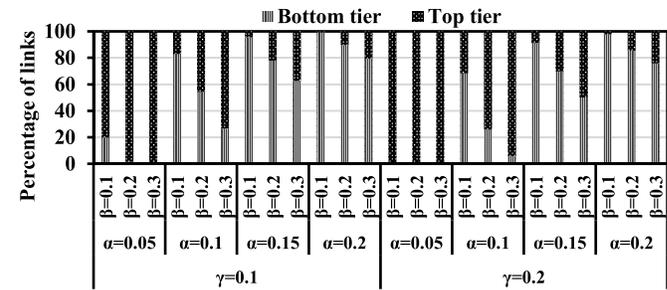


Fig. 7. Tier distribution of links for SWNoC (CAN).

As mentioned in Section IV-C, there is no effect of β on the placement of the crossbar stage (it is not connected directly to an inter-router link). However, the SWA and VCA stages are connected to the router ports, which in turn are connected to the inter-router links. As discussed in Section IV-C, the α and γ parameters affect the intra-router stages and β slows down the inter-router links in the bottom tier. Thus, all three parameters α , β , and γ influence the distribution of the SWA and VCA stages and the inter-router links associated with them. Figs. 6 and 7 show the distribution of the SWA and VCA stages, and inter-router links, respectively, for different values of α , β , and γ for the CAN benchmark. As β increases, the number of BT stages decreases (see Fig. 6) and more links are placed at the top tier (see Fig. 7) to avoid the interconnect performance penalty. For the links in the top tier, the associated stages must become multitier or TT (TT is never chosen since multitier is always better as discussed earlier). So, in Fig. 6, the number of multitier stages increases with the rise of β . On the other contrary, as α increases, multitier stages experience more delay and energy degradation. Hence, more stages (see Fig. 6) and their respective links (see Fig. 7) are placed at the bottom tier to avoid the transistor degradation. For $\gamma = 0.1$, 95.9% of the stages and 97.8% of the links are placed in the bottom tier when we consider the highest value of α ($\alpha = 0.2$) and the lowest value of β ($\beta = 0.1$). Alternatively, all the intra-router stages are multitier, and all the links are placed in the top tier for the lowest value of α ($\alpha = 0.05$) and the highest value of β ($\beta = 0.3$). Overall, the router stages and the inter-router links are distributed to reduce the effects of α and β as much as possible. We can also observe the effect of γ in Figs. 6 and 7. On average (considering different α and β), the number of multitier stages and top-tier links increases by 9.6% and 8%, respectively, when γ increases from 0.1 to 0.2. Hence, considering the effects of various

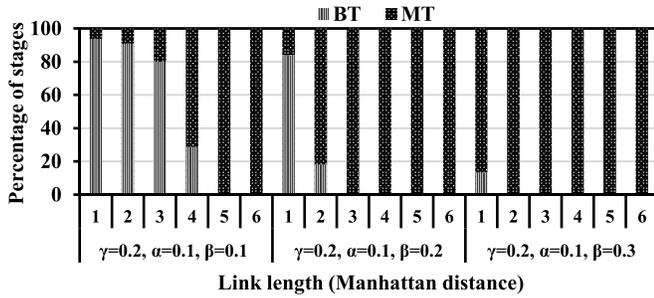


Fig. 8. Tier distribution of the SWA and VCA stages connected to a link of particular length for SWNoC (CAN).

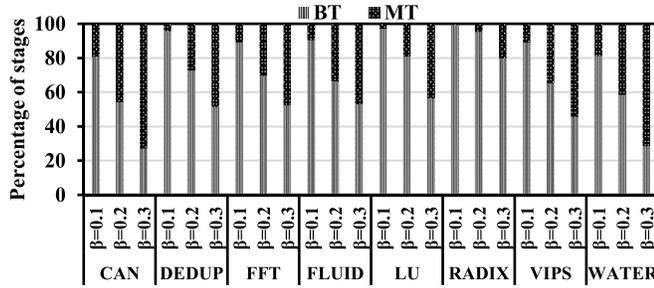


Fig. 9. Tier distribution of the SWA and VCA stages for SWNoC considering all benchmarks ($\alpha = 0.1$, $\gamma = 0.1$).

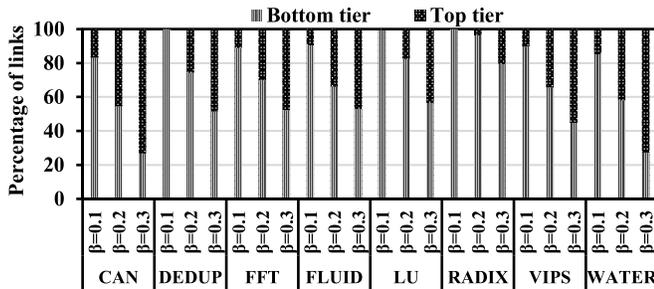


Fig. 10. Tier distribution of links for SWNoC considering all benchmarks ($\alpha = 0.1$, $\gamma = 0.1$).

process variation parameters, the NoC router stages can be placed on a BT or can be distributed over multitiers to optimize EDP.

In Fig. 8, the tier-wise distribution for the stages connected to a link of a particular length is plotted for the CAN benchmark as an example. The link length is expressed in terms of Manhattan distance. As we can see, the placement of SWA and VCA stages is associated with the length of the links. The inter-router traversal penalty at the bottom tier is proportional to the link length (8). Hence, the long-range links are placed mostly in the top tier to avoid the performance penalty (β), whereas the shorter links along with their respective SWA and VCA stages are placed in the bottom tier. Here, the stages connected to longer links favor multitier over BT.

We also found similar trends of the stage and link distribution in all the benchmarks. We plot the input-output stage (SWA and VCA) and link distribution of SWNoCs in Figs. 9 and 10 for representative values $\alpha = 0.1$ and $\gamma = 0.1$ with varying β . As previously discussed, there are zero TT stages and the number of BT stages and bottom-tier

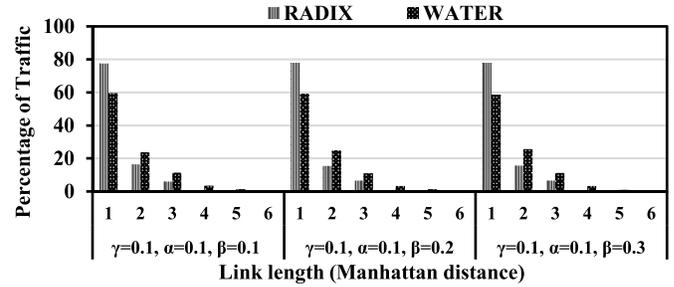


Fig. 11. Percentage of traffic exchanged between any two routers as a function of Manhattan distance for RADIX and WATER.

links decreases with increasing β across all benchmarks. The placement of intra-router stages connected to links (SWA and VCA) and inter-router links varies depending on the traffic characteristics of the specific benchmark. To understand this benchmark dependence, we analyzed RADIX (high percentage of BT and bottom-tier links) and WATER (low percentage of BT and bottom-tier links), two benchmarks that exhibit distinct link distribution trends. In Fig. 11, we show the percentage of traffic exchanged between two routers separated by a particular Manhattan distance. For two routers separated by one Manhattan distance, the traffic exchanged is significantly more for RADIX compared to WATER (18.6% on average). Moreover, RADIX has almost no traffic between routers separated by a Manhattan distance greater than three. Since traffic in RADIX does not have to travel physically far, much of the RADIX traffic takes short-distance links. As we found in the analysis of Fig. 8, this influences the network to have short bottom-tier interconnects. On the other hand, WATER has more traffic that needs to travel further, causing fewer bottom-tier links, especially for higher values of β . In fact, all the intra-router stages are BT and all the inter-router links are placed at the bottom tier for RADIX ($\alpha = 0.1$, $\gamma = 0.1$, and $\beta = 0.1$) as 77.6% of the total traffic is exchanged between routers separated by one Manhattan distance (see Fig. 11). Hence, the link and stage distribution depend on the degree of process variation and traffic pattern of the respective benchmark.

F. Optimization of Mesh-Based NoC

So far, we have considered the SWNoC architecture to thoroughly study the effects of process variation on the NoC router configuration. To study the effects of the process variation on a regular NoC architecture, we undertake the same analysis on an equivalent mesh NoC of the same size. Although the crossbar stage distribution of mesh NoCs is similar to that of SWNoCs, the input-output stages' and links' tier placement are different. We show the distribution of stages (SWA and VCA) and links for a mesh NoC in Figs. 12 and 13, respectively, considering the CAN benchmark. Compared to the stage (see Fig. 6) and link (see Fig. 7) distributions in SWNoC, the SWA and VCA stages, and links in mesh NoC favor the bottom tier. This is attributed to the fact that a mesh NoC consists of only short-range links between adjacent routers causing the router logic to have greater influence on delay and energy. Therefore, in order to avoid the top-tier

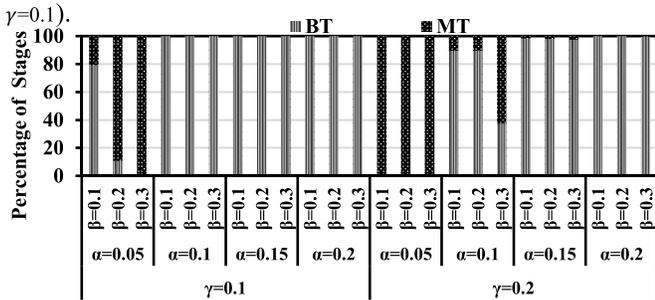


Fig. 12. Tier distribution of the SWA and VCA stages for mesh NoC (CAN).

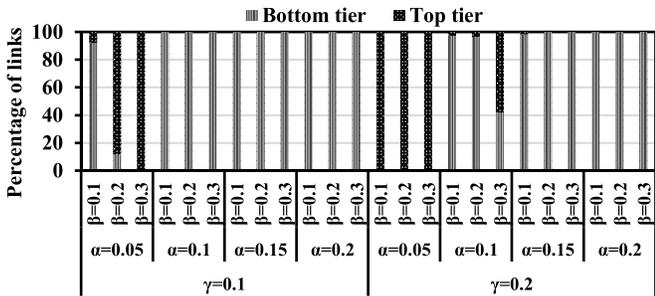


Fig. 13. Tier distribution of links for mesh NoC (CAN).

transistor penalty (except for low values of α), links are placed in the bottom tier. On the other hand, SWNoCs contain both short- and long-range links. The long-range links placed between nonadjacent routers incur more delay and energy overhead if they are placed in the bottom tier.

The general characteristics of process-aware design in SWNoCs are also present in their mesh counterparts. As γ increases, the number of multitier stages also increases as seen in Fig. 12. Similarly, the number of BT stages (see Fig. 12) and bottom-tier links (see Fig. 13) increase with lower values of β or higher values of α . Hence, our process-aware design accounts for the effects of α and β for both SWNoC and mesh NoC.

G. Process-Oblivious Versus Process-Aware M3D NoC

In this section, our aim is to demonstrate the advantages of designing the M3D NoC when considering the effects of process variation. As explained above, when we consider the effects of process variation, the intra-router stages and inter-router links need to be distributed suitably among the tiers. We call this M3D NoC the process-aware architecture (M3D-PA). On the contrary, if we do not consider the effects of process variation (by assuming $\alpha = 0$, $\beta = 0$) while designing the M3D NoC, every router would be multitier to take advantage of the performance benefits due to γ , we call this M3D NoC the process-oblivious M3D NoC (M3D-PO). Our aim is to quantify the benefits of our process-aware design compared to its process-oblivious counterpart.

Figs. 14 and 15 show the EDP of the SWNoC and mesh NoC, respectively, for the CAN benchmark. The EDP is normalized with respect to the M3D-PO design under ideal conditions ($\alpha = 0$, $\beta = 0$). For SWNoCs at the lowest value of α ($\alpha = 0.05$), the M3D-PA design improves the EDP by 10.7% and 9.1% on average considering all values of β over its

M3D-PO counterpart for $\gamma = 0.1$ and $\gamma = 0.2$, respectively. Since the stage distribution of M3D-PO and M3D-PA design is similar for $\alpha = 0.05$, the difference in EDP between these two design approaches is low. In addition, at low values of α most of the stages are multitier (see Fig. 4), this allows the optimization to utilize the top layer for links, reducing the effects of beta on EDP.

As the value of α increases, the multitier router stages get increasingly penalized, so the M3D-PA designs use fewer multitier and more BT stages as shown in Fig. 4. Therefore, there is more opportunity to make better decisions by establishing the trade-off between α and β . However, since the M3D-PO designs do not consider the process variation, the EDP difference between the M3D-PA and M3D-PO designs becomes more prominent. For the most severe process variation ($\alpha = 0.2$, $\beta = 0.3$), the M3D-PA SWNoCs outperform the M3D-PO counterparts by 43.9% and 37.6% for $\gamma = 0.1$ [see Fig. 14(a)] and $\gamma = 0.2$ [see Fig. 14(b)].

We observe similar trends in the EDP distribution of the mesh NoCs in Fig. 15. In fact, we save more EDP for the M3D-PA mesh NoCs with respect to the M3D-PO counterpart than the SWNoC. For the maximum process variation, the process-aware designs outperform the process-oblivious counterpart by 69.6% and 64.2%, for $\gamma = 0.1$ [see Fig. 15(a)] and $\gamma = 0.2$ [see Fig. 15(b)]. In the mesh NoCs, we need more hops to traverse between a pair of source and destination routers on average. This results in passing through more intra-router stages and inter-router links which relate to more opportunities to make the appropriate trade-offs due to the process variation parameters (α and β).

We show the EDP of all benchmarks in Fig. 16 for the SWNoC and the mesh NoC. We chose three different process variation combinations that cover the range of possible values: $\alpha = 0.1$, $\beta = 0.1$ (LOW), $\alpha = 0.15$, $\beta = 0.2$ (MED), and $\alpha = 0.2$, $\beta = 0.3$ (HIGH) to observe the effects of different levels of process variations while keeping γ at 0.1. As discussed above, the EDP of the M3D-PO design deteriorates as the value of α and β increases. For the SWNoCs [see Fig. 16(a)], on average, the M3D-PA saves 19.6%, 33.1%, and 48.7% EDP compared to its M3D-PO counterpart for LOW, MED, and HIGH. For the mesh NoCs [see Fig. 16(b)], the M3D-PA design saves 27.5%, 47.9%, and 70.2% EDP on average compared to its M3D-PO counterpart for LOW, MED, and HIGH, respectively.

H. EDP Comparison Between TSV- and M3D-Based SWNoCs

To complete the analysis, we compare the M3D-PA SWNoC with respect to the TSV-based counterpart of the same size [38]. In Fig. 17, we present the EDP of TSV and M3D-PA SWNoCs. The EDP is normalized with respect to the TSV-based design. Here, we consider the maximum effect of process variation ($\alpha = 0.2$, $\beta = 0.3$) and the lowest value of γ ($\gamma = 0.1$). It is evident from the figure that even in the worst case for process variations, M3D-based designs still reduce the EDP by 11.7% on average compared to TSV based designs for all benchmarks.

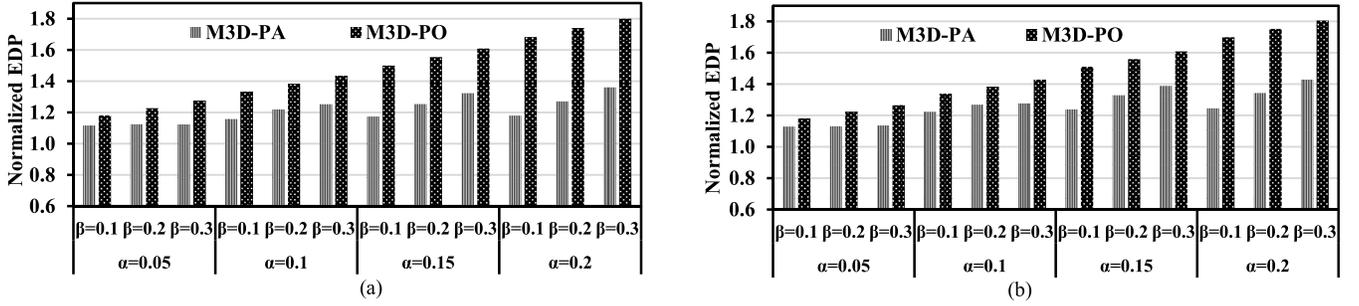


Fig. 14. EDP for SWNoCs for (a) $\gamma = 0.1$ and (b) $\gamma = 0.2$ (CAN). EDP is normalized with respect to that of process-oblivious design under ideal conditions ($\alpha = 0, \beta = 0$).

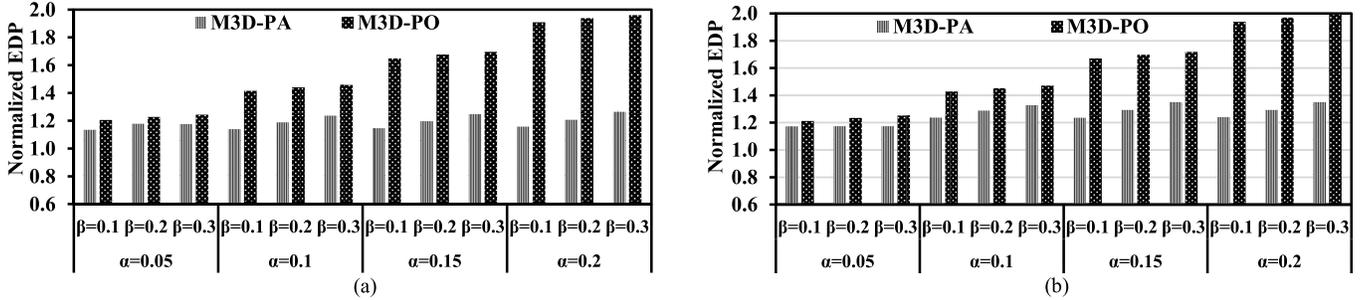


Fig. 15. EDP for mesh NoCs for (a) $\gamma = 0.1$ and (b) $\gamma = 0.2$ (CAN). EDP is normalized with respect to that of process-oblivious design under ideal conditions ($\alpha = 0, \beta = 0$).

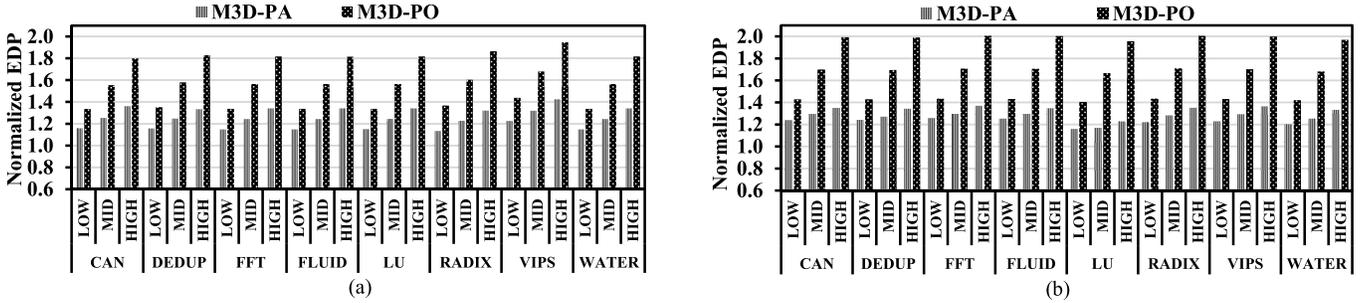


Fig. 16. EDP for (a) SWNoCs and (b) mesh NoCs considering all benchmarks ($\gamma = 0.1$). EDP is normalized with respect to that of process-oblivious design under ideal conditions ($\alpha = 0, \beta = 0$).

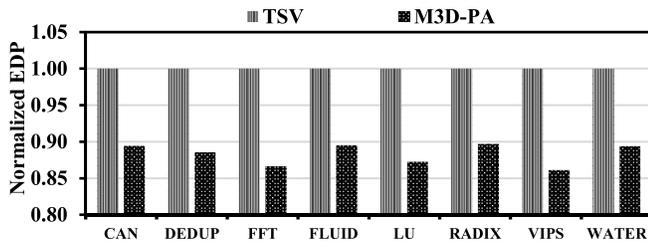


Fig. 17. EDP of TSV- and M3D-enabled ($\gamma = 0.1, \alpha = 0.2, \beta = 0.3$) SWNoCs. EDP is normalized with respect to that of the TSV-based design.

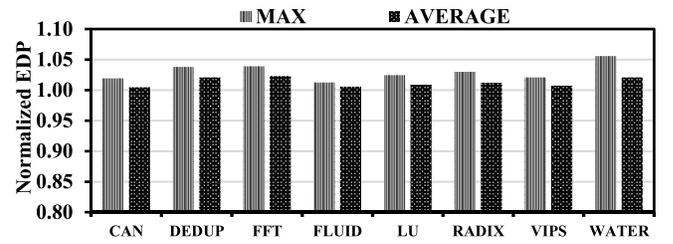


Fig. 18. Normalized EDP for running different applications on the process-aware M3D SWNoC optimized for average f_{ij} .

I. Application Agnostic Process-Aware M3D NoC Design

So far, we have considered the NoC design to be optimized for each benchmark. However, we can design an application agnostic NoC with a simple modification to the definition of f_{ij} (communication traffic). If we want to run multiple applications on the same NoC design, we would consider the average traffic pattern of all the benchmarks and then design and optimize the NoC using this averaged traffic pattern.

To be more specific, suppose we have M benchmark with different inter-core communication traffic patterns, f_{ijk} , where $k = 1, 2, \dots, M$. We then consider the average f_{ijk} in the optimization process and create a single NoC for all benchmarks. In Fig. 18, we present the normalized EDP of the process-aware SWNoC (called SWNoC-A) optimized using the averaged f_{ij} with respect to the process-aware SWNoC (called SWNoC-S) optimized using each specific application's f_{ij} . For example, CAN shows the EDP of the SWNoC

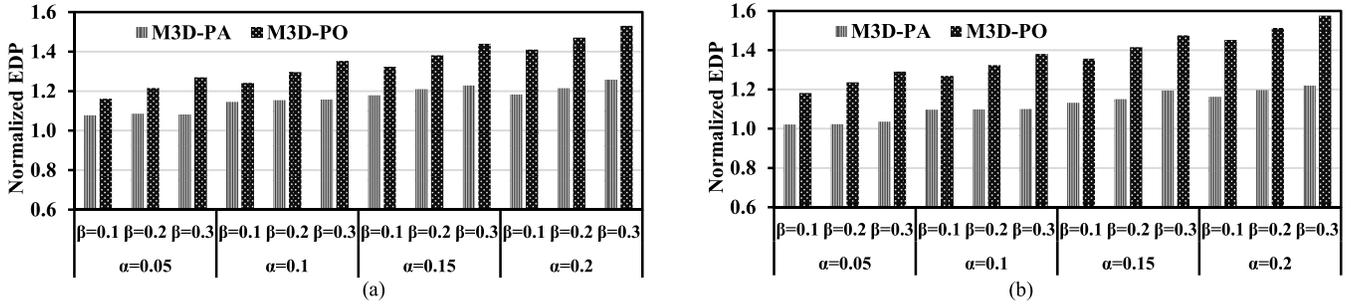


Fig. 19. EDP for SWNoCs at (a) $\gamma = 0.1$ and (b) $\gamma = 0.2$ for FFT benchmark (MCSL suite) and system size of 128. EDP is normalized with respect to that of process-oblivious design under ideal conditions ($\alpha = 0, \beta = 0$).

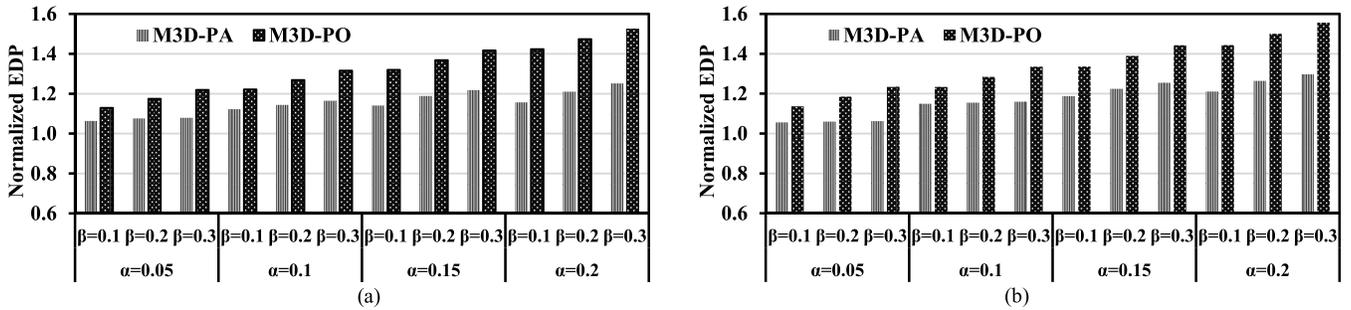


Fig. 20. EDP for SWNoCs at (a) $\gamma = 0.1$ and (b) $\gamma = 0.2$ for FFT benchmark (MCSL suite) and system size of 256. EDP is normalized with respect to that of process-oblivious design under ideal conditions ($\alpha = 0, \beta = 0$).

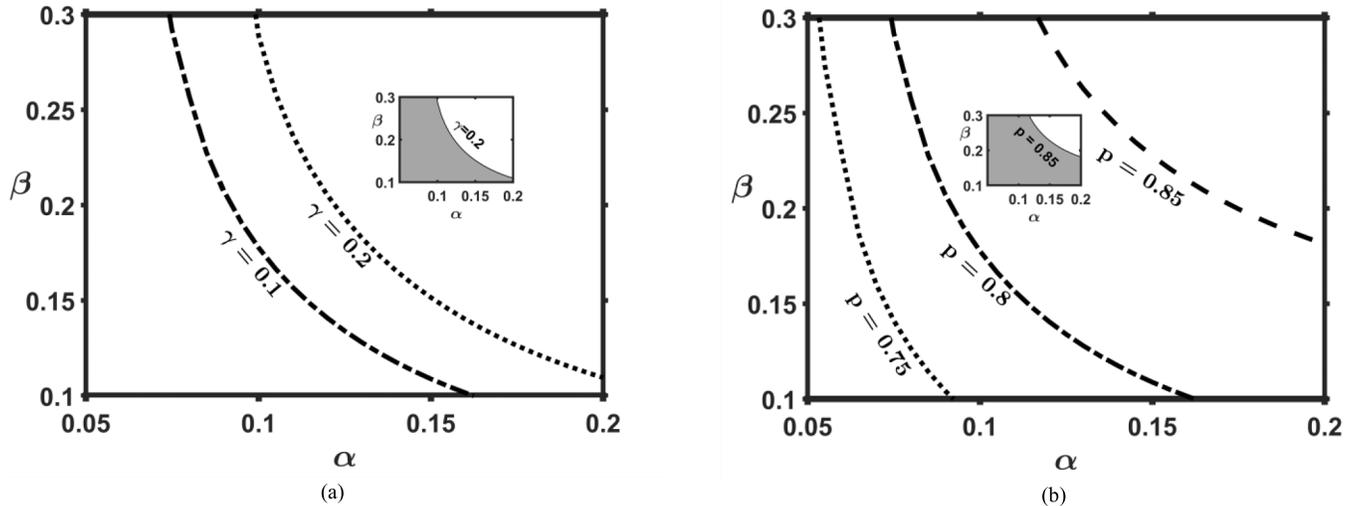


Fig. 21. Largest combinations of α and β that manage to satisfy (9) for (a) different values of γ for $p = 0.8$ and (b) different values p for $\gamma = 0.1$. All values of α and β that lie under the curve satisfy (9) and the shaded area within the inset graphs demonstrate two examples. CAN is considered here as an example.

optimized using the averaged f_{ij} normalized to the SWNoC optimized using CAN's f_{ij} . Considering all values of α , β , and γ , we present the worst case (max) and average EDP for each benchmark. Obviously, the EDP of the SWNoC-A is greater than that of the SWNoC-S because SWNoC-A is optimized using the average traffic pattern. However, the worst case degradation of SWNoC-A is only 5.6% (WATER benchmark) and the average worst case degradation of SWNoC-A is 3% compared to SWNoC-S. Here, we show the results only for SWNoCs for brevity. The performance degradation is insignificant for mesh NoCs. This demonstrates that a general

NoC can be created specifically for PARSEC and SPLASH-2 benchmark suites that run any application without noticeable performance loss.

J. System Size Scalability

In this section, we compare the performance of M3D-PA and M3D-PO SWNoCs for higher system sizes ($N = 128$ and 256). To consider larger system sizes, we consider the MCSL benchmark suite since it extracts realistic traffic patterns of real applications by executing a cycle-accurate simulator. Specifically, we examine the

FFT benchmark from the MCSL suite [45] for NoC design and performance evaluation. To undertake the performance evaluation, we used an in-house cycle-accurate NoC simulator [46]. We are unable to undertake a full system simulation using GEM5 as it does not support systems with more than 64 cores. We show the EDP of M3D-PA and M3D-PO designs for systems with 128 cores in Fig. 19 and 256 cores in Fig. 20. The EDP is normalized with respect to the M3D-PO design under ideal conditions ($\alpha = 0, \beta = 0$). On average, considering different values of α , β , and γ , M3D-PA outperforms the M3D-PO counterpart by 21.5% and 16.8% for system sizes of 128 and 256, respectively.

K. M3D Process Parameter Guidelines for M3D NoCs

Using our analytical formulation to quickly find optimal solutions, we can determine the region for the α , β , and γ parameters that will make the process-aware design more efficient than its TSV-based counterpart. This is important for understanding the design tradeoffs and the inflection point to using M3D over TSV-based designs. To achieve that, we need to find the values of α , β , and γ that satisfy the following:

$$\text{EDP}_{\text{M3D}} \leq p \cdot \text{EDP}_{\text{TSV}} \quad (9)$$

where EDP_{M3D} and EDP_{TSV} are the EDPs of the M3D- and TSV-based NoCs, respectively, and p represents the target EDP ratio between them. Then, we do an exhaustive sweep over the ranges of α , β , and γ given in Table III and create process-aware designs for each step. In Fig. 21, we demonstrate the upper bounds of α and β that the process-aware design can satisfy the equality constraint in (9) under certain conditions. In Fig. 21(a), at an aggressive target $p = 0.8$, we show two different upper bound curves for $\gamma = 0.1$ and $\gamma = 0.2$ considering the CAN benchmark. The corresponding area under the curve (for the $\gamma = 0.2$ curve, the area is shaded in the inset plot) represents the range of α and β that our process-aware design can tolerate and still achieve at least 20% performance benefits compared to a TSV-based design. As the value of γ increases (larger benefits for using M3D), more α and β values can be tolerated while satisfying (9). Fig. 21(b) shows a similar plot at $\gamma = 0.1$ for various values of p . As we decrease p , the curves move closer to the origin and reduces the acceptable ranges of α and β . We see these trends for all benchmarks and other values of γ . We do not repeat those results here for brevity.

VI. CONCLUSION

Although M3D-integration offers high performance and energy efficiency, fabrication-related challenges pose major concerns to achieve desirable performance levels. The process-induced performance degradation of the transistors and interconnects introduces significant performance and energy overheads for M3D-enabled NoCs. Our analysis shows that the SWNoC designed without considering the process variation underestimates the EDP by at least 18.8% and at most 83.7% depending on the process variation parameters for 64-core based system. Thus, process-variation-aware design is a must

for realistic M3D NoC architectures. In this article, we incorporated both top-tier transistor slow-down and bottom-tier interconnect degradation in the NoC design process. Our proposed design reduces the performance degradation of the M3D NoC by suitably distributing the intra-router stages and inter-router links among the M3D tiers. The process-aware design improves the EDP of SWNoC by at least 7.2% and up to 48.7% compared to the process-oblivious design approach for the best and worst case of process variation (for 64-core based system), respectively. Most importantly, although the natural impulse is to make the entire system multitier, all routers should not be made multitier. We demonstrated that this is the case even for different NoC topologies and architectures. Depending on the process variation and traffic patterns, various parts of the NoC routers should be placed in different tiers.

REFERENCES

- [1] W. R. Davis *et al.*, "Demystifying 3D ICs: The pros and cons of going vertical," *IEEE Design Test Comput.*, vol. 22, no. 6, pp. 498–510, Nov./Dec. 2005.
- [2] B. S. Feero and P. P. Pande, "Networks-on-chip in a three-dimensional environment: A performance evaluation," *IEEE Trans. Comput.*, vol. 58, no. 1, pp. 32–45, Jan. 2009.
- [3] X. Dong and Y. Xie, "System-level cost analysis and design exploration for three-dimensional integrated circuits (3D ICs)," in *Proc. Asia South Pacific Design Autom. Conf.*, Yokohama, Japan, 2009, pp. 234–241.
- [4] J. H. Lau, "TSV manufacturing yield and hidden costs for 3D IC integration," in *Proc. 60th Electron. Compon. Technol. Conf. (ECTC)*, Las Vegas, NV, USA, Jun. 2010, pp. 1031–1042.
- [5] K. Athikulwongse, A. Chakraborty, J.-S. Yang, D. Z. Pan, and S. K. Lim, "Stress-driven 3D-IC placement with TSV keep-out zone and regularity study," in *Proc. IEEE/ACM Int. Conf. Comput.-Aided Design (ICCAD)*, San Jose, CA, USA, Nov. 2010, pp. 669–674.
- [6] S. Das, J. R. Doppa, P. P. Pande, and K. Chakrabarty, "Robust TSV-based 3D NoC design to counteract electromigration and crosstalk noise," in *Proc. Design, Autom. Test Eur. Conf. Exhib.*, Lausanne, Switzerland, Mar. 2017, pp. 1366–1371.
- [7] P. Batude, T. Ernst, J. Arcamone, G. Arndt, P. Coudrain, and P. E. Gaillardon, "3-D sequential integration: A key enabling technology for heterogeneous co-integration of new function with CMOS," *IEEE J. Emerg. Sel. Topics Circuits Syst.*, vol. 2, no. 4, pp. 714–722, Dec. 2012.
- [8] S. K. Samal, D. Nayak, M. Ichihashi, S. Banna, and S. K. Lim, "Monolithic 3D IC vs. TSV-based 3D IC in 14nm FinFET technology," in *Proc. IEEE SOI-3D-Subthreshold Microelectron. Technol. Unified Conf.*, Burlingame, CA, USA, Oct. 2016, pp. 1–2.
- [9] S.-M. Jung *et al.*, "High speed and highly cost effective 72M bit density S³ SRAM technology with doubly stacked Si layers, peripheral only CoSix layers and tungsten shunt W/L scheme for standalone and embedded memory," in *Proc. IEEE Symp. VLSI Technol.*, Kyoto, Japan, Jun. 2007, pp. 82–83.
- [10] T. Uhrmann, T. Wagenleitner, T. Glinsner, M. Wimplinger, and P. Lindner, "Monolithic IC integration key alignment aspects for high process yield," in *Proc. SOI-3D-Subthreshold Microelectron. Technol. Unified Conf.*, Millbrae, CA, USA, Oct. 2014, pp. 1–2.
- [11] S. Das, J. R. Doppa, P. P. Pande, and K. Chakrabarty, "Monolithic 3D-enabled high performance and energy efficient network-on-chip," in *Proc. IEEE Int. Conf. Comput. Design (ICCD)*, Boston, MA, USA, Nov. 2017, pp. 233–240.
- [12] S.-E. D. Lin and D. H. Kim, "Detailed-placement-enabled dynamic power optimization of multitier gate-level monolithic 3-D ICs," *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, vol. 37, no. 4, pp. 845–854, Apr. 2018.
- [13] B. Rajendran *et al.*, "Low thermal budget processing for sequential 3-D IC fabrication," *IEEE Trans. Electron Devices*, vol. 54, no. 4, pp. 707–714, Apr. 2007.
- [14] M. Lin, A. El Gamal, Y.-C. Lu, and S. Wong, "Performance benefits of monolithically stacked 3-D FPGA," *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, vol. 26, no. 2, pp. 216–229, Feb. 2007.

- [15] S. Bobba, A. Chakraborty, O. Thomas, P. Batude, and G. de Micheli, "Cell transformations and physical design techniques for 3D monolithic integrated circuits," *ACM J. Emerg. Technol. Comput. Syst.*, vol. 9, no. 3, p. 19, Sep. 2013.
- [16] Y.-J. Lee, P. Morrow, and S. K. Lim, "Ultra high density logic designs using transistor-level monolithic 3D integration," in *Proc. IEEE/ACM Int. Conf. Comput.-Aided Design (ICCAD)*, San Jose, CA, USA, Nov. 2012, pp. 539–546.
- [17] K. Acharya *et al.*, "Monolithic 3D IC design: Power, performance, and area impact at 7nm," in *Proc. 17th Int. Symp. Qual. Electron. Design (ISQED)*, Santa Clara, CA, USA, Mar. 2016, pp. 41–48.
- [18] K. M. Kim, S. Sinha, B. Cline, G. Yeric, and S. K. Lim, "Four-tier monolithic 3D ICs: Tier partitioning methodology and power benefit study," in *Proc. Int. Symp. Low Power Electron. Design*, San Francisco, CA, USA, Aug. 2016, pp. 70–75.
- [19] C. Yan and E. Salman, "Mono3D: Open source cell library for monolithic 3-D integrated circuits," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 65, no. 3, pp. 1075–1085, Mar. 2018.
- [20] B. Gopireddy and J. Torrellas, "Designing vertical processors in monolithic 3D," in *Proc. 46th Int. Symp. Comput. Archit.*, Phoenix, AZ, USA, Jun. 2019, pp. 643–656.
- [21] M. M. S. Aly *et al.*, "Energy-efficient abundant-data computing: The N3XT 1,000x," *Computer*, vol. 48, no. 12, pp. 24–33, Dec. 2015.
- [22] A. Ayres *et al.*, "Variance analysis in 3-D integration: A statistically unified model with distance correlations," *IEEE Trans. Electron Devices*, vol. 66, no. 1, pp. 633–640, Jan. 2019.
- [23] A. Tsiara *et al.*, "Performance and reliability of a fully integrated 3D sequential technology," in *Proc. IEEE Symp. VLSI Technol.*, Honolulu, HI, USA, Jun. 2018, pp. 75–76.
- [24] A. Koneru, S. Kannan, and K. Chakrabarty, "Impact of electrostatic coupling and wafer-bonding defects on delay testing of monolithic 3D integrated circuits," *ACM J. Emerg. Technol. Comput. Syst.*, vol. 13, no. 4, p. 54, Aug. 2017.
- [25] T. Naito *et al.*, "World's first monolithic 3D-FPGA with TFT SRAM over 90nm 9 layer Cu CMOS," in *Proc. Symp. VLSI Technol.*, Honolulu, HI, USA, Jun. 2010, pp. 219–220.
- [26] N. K. Jha and Y. Yu, "Energy-efficient monolithic three-dimensional on-chip memory architectures," *IEEE Trans. Nanotechnol.*, vol. 17, no. 4, pp. 620–633, Jul. 2018.
- [27] N. Golshani *et al.*, "Monolithic 3D integration of SRAM and image sensor using two layers of single grain silicon," in *Proc. IEEE Int. 3D Syst. Integr. Conf. (3DIC)*, Munich, Germany, Nov. 2010, pp. 1–4.
- [28] D. Lee, S. Das, J. R. Doppa, P. P. Pande, and K. Chakrabarty, "Performance and thermal tradeoffs for energy-efficient monolithic 3D network-on-chip," *ACM Trans. Design Autom. Electron. Syst.*, vol. 23, no. 5, p. 60, Oct. 2018.
- [29] S. Panth, S. K. Samal, K. Samadi, Y. Du, and S. K. Lim, "Tier degradation of monolithic 3-D ICs: A power performance study at different technology nodes," *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, vol. 36, no. 8, pp. 1265–1273, Aug. 2017.
- [30] L. Pasini *et al.*, "nFET FDSOI activated by low temperature solid phase epitaxial regrowth: Optimization guidelines," in *Proc. SOI-3D-Subthreshold Microelectron. Technol. Unified Conf. (S3S)*, Millbrae, CA, USA, Oct. 2014, pp. 1–2.
- [31] C. Fenouillet-Beranger *et al.*, "New insights on bottom layer thermal stability and laser annealing promises for high performance 3D VLSI," in *Proc. IEEE Int. Electron Device Meeting*, San Francisco, CA, USA, Dec. 2014, pp. 27.5.1–27.5.4.
- [32] L.-S. Peh and W. J. Dally, "A delay model for router microarchitectures," *IEEE Micro*, vol. 21, no. 1, pp. 26–34, Jan. 2001.
- [33] H.-S. Wang, L.-S. Peh, and S. Malik, "A power model for routers: Modeling Alpha 21364 and InfiniBand routers," in *Proc. 10th Symp. High Perform. Interconnects*, Stanford, CA, USA, Aug. 2002, pp. 21–27.
- [34] A. Mallik *et al.*, "The impact of sequential-3D integration on semiconductor scaling roadmap," in *IEDM Tech. Dig.*, San Francisco, CA, USA, Dec. 2017, pp. 32.1.1–31.1.4.
- [35] P. Batude *et al.*, "3D monolithic integration," in *Proc. IEEE Int. Symp. Circuits Syst. (ISCAS)*, Rio de Janeiro, Brazil, May 2011, pp. 2233–2236.
- [36] S. K. Samal, D. Nayak, M. Ichihashi, S. Banna, and S. K. Lim, "Tier partitioning strategy to mitigate BEOL degradation and cost issues in monolithic 3D ICs," in *Proc. IEEE/ACM Int. Conf. Comput.-Aided Design (ICCAD)*, Austin, TX, USA, Nov. 2016, pp. 1–7.
- [37] J. Boyan and A. W. Moore, "Learning evaluation functions to improve optimization by local search," *J. Mach. Learn. Res.*, vol. 1, pp. 77–112, Jan. 2000.
- [38] S. Das, J. R. Doppa, P. P. Pande, and K. Chakrabarty, "Design-space exploration and optimization of an energy-efficient and reliable 3-D small-world network-on-chip," *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, vol. 36, no. 5, pp. 719–732, May 2017.
- [39] N. Binkert *et al.*, "The gem5 simulator," *ACM SIGARCH Comput. Archit. News*, vol. 39, no. 2, pp. 1–7, 2011.
- [40] S. C. Woo, M. Ohara, E. Torrie, J. P. Singh, and A. Gupta, "The SPLASH-2 programs: Characterization and methodological considerations," in *Proc. 22nd Annu. Int. Symp. Comput. Archit.*, Santa Margherita Ligure, Italy, Jun. 1995, pp. 24–36.
- [41] C. Bienia, "Benchmarking modern multiprocessors," Ph.D. dissertation, Dept. Comput. Sci., Princeton Univ., Princeton, NJ, USA, 2011.
- [42] N. Barrow-Williams, C. Fensch, and S. Moore, "A communication characterisation of Splash-2 and Parsec," in *Proc. IEEE Int. Symp. Workload Characterization (IISWC)*, Austin, TX, USA, Oct. 2009, pp. 86–97.
- [43] O. Lysne, T. Skeie, S.-A. Reinemo, and I. Theiss, "Layered routing in irregular networks," *IEEE Trans. Parallel Distrib. Syst.*, vol. 17, no. 1, pp. 51–65, Jan. 2006.
- [44] K. Chang *et al.*, "Cascade2D: A design-aware partitioning approach to monolithic 3D IC with 2D commercial tools," in *Proc. IEEE/ACM Int. Conf. Comput.-Aided Design (ICCAD)*, Austin, TX, USA, Nov. 2016, pp. 1–8.
- [45] W. Liu *et al.*, "A NoC traffic suite based on real applications," in *Proc. IEEE Comput. Soc. Annu. Symp. VLSI*, Chennai, India, Jul. 2011, pp. 66–71.
- [46] P. Wettin *et al.*, "Design space exploration for wireless NoCs incorporating irregular network routing," *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.*, vol. 33, no. 11, pp. 1732–1745, Nov. 2014.