ARTEMIS: An Aging-Aware Runtime Application Mapping Framework for 3D NoC-Based Chip Multiprocessors

Venkata Yaswanth Raparti, *Student Member, IEEE*, Nishit Kapadia, *Member, IEEE*, and Sudeep Pasricha, *Senior Member, IEEE*

Abstract—In emerging 3D NoC-based chip multiprocessors (CMPs), aging in circuits due to bias temperature instability (BTI) stress is expected to cause gate-delay degradation that, if left unchecked, can lead to untimely failure. Simultaneously, the effects of electromigration (EM) induced aging in the on-chip wires, especially those in the 3D power delivery network (PDN), are expected to notably reduce chip lifetime. A commonly proposed solution to mitigate circuit-slowdown due to aging is to hike the supply voltage; however, this increases current-densities in the PDN due to the increased power consumption on the die, which in turn expedites PDN-aging. We thus note that mechanisms to enhance lifetime reliability in 3D NoC-based CMPs must consider circuit-aging together with PDN-aging. In this paper, we propose a novel runtime framework (*ARTEMIS*) for intelligent dynamic application-mapping and voltage-scaling to simultaneously manage aging in circuits and the PDN, and enhance the performance and lifetime of 3D NoC-based CMPs. We also propose an aging-enabled routing algorithm that balances the degree of aging between NoC routers and cores, thereby increasing the combined lifetime of both. Our framework also considers dark-silicon power constraints that are becoming a major design challenge in scaled technologies, particularly for 3D stacked CMPs. Our experimental results indicate that *ARTEMIS* enables the execution of 25 percent more applications over the chip lifetime compared to state-of-the-art prior work.

Index Terms—Application mapping, aging resilience, power delivery network, dark-silicon, 3D chip multiprocessors

1 Introduction

Bias Temperature Instability (BTI) is the most dominant physical phenomenon that degrades the maximum switching rate of transistors under long periods of voltage stress in emerging chip multiprocessors (CMPs) [1]. BTI causes gradual circuit slowdown over the operational lifetime of the electronic chip. For systems manufactured at technology nodes below 45 nm, BTI-induced delay-degradation can be quite significant [2], [3]. The principal effect of such a circuit-aging mechanism is to increase circuit-threshold voltage (V_T), which results in higher circuit-delay. From a system-level perspective, such V_T-degradation causes slowdown in critical paths of processor-cores and network-on-chip (NoC) routers, thereby limiting overall system performance. With increasing demand for reliable CMPs with longer lifetimes in non-consumer domains such as aerospace, defense, automobile, and health, prolonging the useful CMP lifetime will be very beneficial.

- *V.Y. Raparti and S. Pasricha are with the Department of Electrical and Computer Engineering, Colorado State University, Fort Collins, CO 80523. E-mail: yaswanth@rams.colostate.edu, sudeep@colostate.edu.*
- *N. Kapadia is with Synopsys Inc., Hillsboro, OR 97124. E-mail: nishit.kapadia@synopsys.com.*

Manuscript received 24 Aug. 2016; revised 9 Mar. 2017; accepted 9 Mar. 2017. Date of publication 23 Mar. 2017; date of current version 15 June 2017. Recommended for acceptance by K. Chakrabarty.
For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org, and reference the Digital Object Identifier below.
Digital Object Identifier no. 10.1109/TMSCS.2017.2686856

Additionally, electromigration (EM) in metal wires on the chip leads to increased interconnect resistance over time in CMPs. This phenomenon is most dominant in power delivery network (PDN) wires that carry larger unidirectional currents compared to signal wires [4], [12]. The increased resistance of the power-grid results in higher IR-drops in the PDN, which causes further circuit slowdown due to degradation of supply voltage [5]. These adverse effects of EM are expected to be particularly severe in 3D CMPs that possess limited number of power-pins and higher current densities [6]. Also, with process technology scaling, this problem is exacerbated due to the reduction in cross-sections of metal wires, which causes further increase in PDN current-densities [4].

To mitigate BTI-induced delay degradation, while maintaining circuit operation at a minimum clock frequency (i.e., minimum performance level), one solution is to hike the supply voltage [13] adaptively over time based on the degree of circuit-aging. However, doing so increases current-densities in the PDN due to the increased power dissipated in the chip as a result of the voltage-hike. High current densities end up causing faster EM-induced PDN-aging [7], hastening circuit-slowdown. Hiking supply voltage also increases V_T-degradation, further increasing the rate of circuit-aging.

As aging reduces the viable lifetime of current and emerging CMPs, it is becoming increasingly important to consider it during the design process. Unfortunately, designers today

are more focused on meeting performance requirements, and resort to either using costly hardware guardbands to minimize the effect of performance variations on a die, or employing large supply voltage guardbands to ensure a reliable voltage supply, that ends up increasing power densities and peak temperature of chip, resulting in shortening chip lifetime. Practical and low-cost solutions to enhance lifetime are thus becoming essential, especially in dense 3D CMPs fabricated in scaled technologies. As noted earlier, such solutions must also consider the interdependence between BTI-induced circuit aging, supply voltage, and EM-induced PDN-aging.

Yet another challenge facing CMP designers is the rise in on-chip power dissipation. The slowdown of power scaling with technology scaling, due to leakage and reliability concerns [8], [9], has led to high chip power-densities, giving rise to the dark-silicon phenomenon, whereby a non-negligible fraction of the chip must be shut down at any given time to satisfy the chip power-budget. With the extent of dark-silicon increasing with every technology-generation [10], [11], designs are becoming increasingly power-limited rather than area-limited. Therefore, runtime power-saving techniques such as dynamic voltage scaling (DVS) are of paramount importance to extract much needed performance given a stringent chip-wide power-budget.

To simultaneously address all the above mentioned challenges related to aging, power dissipation, and performance facing chip designers, in this paper we propose a novel *runtime aging-aware application-mapping framework called ARTEMIS*. Our framework is intended for 3D NoC-based CMPs and aims to increase the useful work performed over the lifetime of these chips, while meeting the dark-silicon power-budget (DS-PB) and application performance goals. The novel contributions of our work are summarized below:

- We propose a novel runtime application-mapping and DVS-scheduling framework that can adapt to different aging scenarios to extend the lifetime of a 3D NoC-based CMP;
- As the impacts of PDN-aging and circuit-aging (for cores and NoC routers) on system-performance are correlated, our framework considers aging in these components, unlike any prior work, while making mapping decisions to alleviate system aging;
- Our methodology to evaluate system-aging and the resulting maximum-attainable performance accounts for progressive effects of IR-drops due to PDN-aging, V_T-degradation due to circuit-aging, and temperature profiles over the chip lifetime;
- We design a novel symmetric aging-enabled NoC routing path allocation (SAR) heuristic to produce a balanced core-router aging profile to extend the lifetime of the NoC. In addition, SAR efficiently trades-off aging with network-congestion in the NoC;
- Our framework also meets chip-wide power constraints, thus finding applicability in contemporary power-constrained (dark-silicon afflicted) multicore designs.

2 Related Work

In recent years, several researchers have proposed run-time and design-time application mapping techniques to address the problem of circuit-aging in CMPs. Tiwari et al. [13] suggest mapping high-power tasks onto faster (less-aged) cores and low-power tasks onto slower cores, thereby "hiding" the aging in the chip. At the same time, they propose to lessen aging by scaling the supply voltage or the threshold voltage. Feng et al. [14] perform "local wear-leveling" by scheduling tasks on cores while considering circuit-aging in sub-core components. But such wear-leveling approaches where "younger" cores are prioritized over aged cores without considering application-performance (frequency) requirements lead to higher leakage power dissipation as faster cores are also leakier, which expedites aging. *Thus, in the dark-silicon regime where performance is closely tied to power, wear-leveling techniques are usually sub-optimal.*

Some of the recent works propose aging-aware frameworks that discretize target lifetime of the chip into finer lifetime constraints, and perform runtime management to satisfy a pre-defined target lifetime and system performance goal. For instance, Mintarno et al. [3] use frequency, voltage, and cooling power as control parameters to optimize the energy-efficiency of a system while meeting lifetime targets. Paterna et al. [15] propose a linear-programming based task-allocation solution to optimize energy, while [16] performs voltage tuning over shorter time-intervals to meet aging constraints over longer time-intervals. But these techniques are either too time consuming to be viable for runtime decision making, or require comprehensive knowledge of future application characteristics, which may not be available in many environments where CMPs are used. In [38], Haghbayan et al. have proposed a lifetime aware runtime mapping framework that maps tasks on to cores to meet the dark-silicon power budget and satisfy the target reliability till the end of the chip lifetime. However, they have not considered other viable resource management approaches such as Dynamic Voltage Scaling (DVS) to execute more applications within a given dark- silicon power budget. They have also considered a 2D CMP where the impact of electro-migration (EM) is not predominant. However, recent chip designers and manufacturers are gravitating towards 3D CMPs, where the degradation (due to EM) in power delivery networks (PDN) can potentially slow down performance of on-chip components such as cores and NoC-routers. In [39], Huang et al. have proposed an analytical model for computing the reliability of a manycore processor, and analyzed the impacts of various redundancy schemes on the lifetime and the performance of the processor. Having redundant cores is a viable solution in a processor with a small core count. But, as the core count increases, redundancy based schemes will incur very high power and area overheads, which is impractical in dark-silicon power constrained chips. In [47], Kapadia et al. have proposed a process variation aware framework for application mapping to improve the reliability due to soft errors in the dark-silicon era.

More recent works have considered on-chip temperature profile as a prime contributor to circuit-aging and propose techniques to reduce the peak on-chip temperature with thermal-aware mapping techniques. Gnad et al., in [40], have proposed a framework that uses an offline generated frequency degradation table for various applications to estimate the aging-induced frequency degradation in cores before mapping, for efficient aging and dark-silicon management

on a 2D CMP. Their framework has objectives similar to *ARTEMIS*, but when applied to a 3D CMP, it ignores PDN-aging. Further, they aim to minimize circuit-aging and power consumption in cores, while ignoring other on-chip components such as NoC. In [41], Singh et al., have proposed a mapping technique to reduce the energy consumption and the peak on-chip temperature while improving the application throughput of 3D video processing applications on 3D CMPs. Even though thermal-aware mapping minimizes aging to an extent in 2D CMPs, mitigating aging in cores and PDN simultaneously in 3D CMPs requires a mapping technique that is cognizant of degradation profiles of both cores and PDN along with the prior knowledge of workload that is executed on them. In [48], [49], Pasricha et al. have proposed a fault tolerant and energy efficient NoC routing scheme for 2D and 3D NoC based systems. In [42], Rehman et al. have proposed a framework to address the reliability challenges due to soft-error induced failures. Their framework aims at leveraging the knowledge of on-chip variation and aging profiles to efficiently choose the hardware-software reliability mitigation and application mapping techniques, to improve the reliability of 2D CMPs. Unlike their framework, we balance circuit- as well as PDN-aging and increase the useful lifetime of 3D CMPs in the presence of dark-silicon. In [43], a novel temperature-model is proposed that considers both spatial and temporal dependencies. This model is used by a design time task scheduling and operation point assignment mechanism, to jointly optimize reliability and energy of multimedia applications represented using streaming data flow graphs (SDFG). However, a design time task assignment does not efficiently capture a run-time aging profile caused by dynamic workload characteristics found in most of today's CMPs. In [44], a run-time task scheduling framework is proposed to minimize the communication energy consumed, without effecting the throughput of multimedia applications on heterogeneous multicore systems in the presence of permanent and intermittent failures. In [45], a scenario-aware fault injection model is proposed to model intermittent failures caused by wear-out mechanisms on 2D MultiProcessor-System-on-Chips (MPSoCs). They further propose a wear-out-aware application mapping technique that maps applications based on the intermittent failure rate, which is an indicator of chip aging. However, this work ignores the presence of dark silicon power constraint that is prevailing in most of today's multicore processors. Also, all these works do not consider the impact of aging in PDN. *To the best of our knowledge, this paper is the first work on lifetime-aware application mapping at runtime that considers the impact of PDN-aging on the lifetime of the chip, and is also tailored for the power-constrained dark-silicon design regime.*

Aging is a concern not just for computation cores but also for NoC fabrics that connect these cores together. But very few works have investigated design-techniques that extend the service life of the NoC fabric. Bhardwaj et al. [35] have proposed an aging-aware adaptive routing algorithm that routes packets along the paths that are both less congested and experience smaller aging stress. But the authors do not consider aging in compute-cores. *Our work represents one of the first efforts to extend the useful lifetime of the entire chip by producing a balanced core-router aging profile, with our proposed symmetric aging-enabled routing path allocation (SAR) heuristic.*

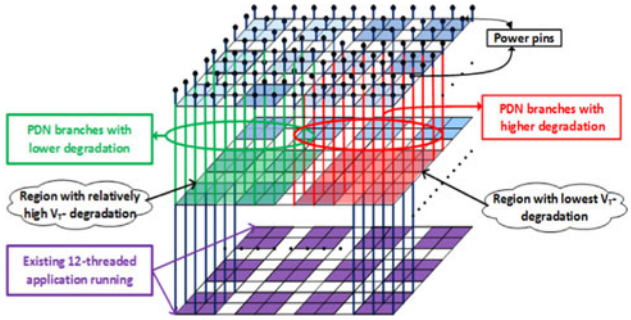

Fig. 1. Example of a 3D package for a 36-core CMP (4x3x3 3D-mesh) with a regular 3D power-grid that has 108 external power-pins and 108 grid-points per tier (16 grid-points supplying to each core). Not all vertical branches of PDN are shown, for brevity.

3 Motivation

In this section, we illustrate the advantages of our *ARTEMIS* framework with the help of a small example. We consider a scenario in which applications arrive at runtime to be executed on a 36-core CMP with a core capable of executing a single thread (task) at a time. The example assumes a 4-thread application being mapped on to the cores of a 3D-CMP at time t, when a 12-thread application is already executing on the bottom tier (shown in purple in Fig. 1). In a 3D-CMP, a 2D region of tiles in a central layer with less V_T-*degradation* can have a relatively high PDN degradation due to the current flowing through the PDN in middle layer, to supply the applications that ran in the bottom layer. The variation in PDN degradation is depicted by red and green colored tiles and lines in Fig. 1 where red PDN lines supplied more current to applications in the bottom layer than green PDN lines. When an aging-aware wear-leveling technique based on prior work [13], [14] is used, the application would be mapped to the rectangular region (shown in red) containing cores with the least V_T-degradation, i.e., the region with the youngest cores. Observe that the vertical PDN branches supplying to this region have high degradation (higher resistance values due to past stress). Alternatively, although the green rectangular region has more V_T-degradation, the PDN IR-drops sustained by it are lower than the red region. By always prioritizing mapping of applications on to the youngest cores, without considering the resulting impact on EM-induced degradation in the PDN, resistances of already stressed PDN-wires would be further increased, thus exacerbating PDN-degradation. Therefore, *ARTEMIS* considers both V_T-degradation and PDN-degradation while making mapping decisions to limit PDN-degradation while guaranteeing that performance and power constraints are met on a 3D NoC-based CMP.

Additionally, the maximum frequency (f_{max}) of a core is affected by both the PDN IR-drops (which affects V_{dd}) as well as V_T-degradation:

$$f_{max} = \frac{\mu(V_{dd} - V_T)^{\alpha}}{C_0.V_{dd}} \quad (1)$$

where α and μ are technology-dependent constants, and C_0 is switching capacitance of the critical path [25]. Approximate values of these constants are listed in . Thus, any mapping solution obtained without consideration of IR-drops experienced by cores can potentially lead to undesirable timing-errors.

In summary, the wear-leveling based application-mapping approach (i.e., always choosing the youngest cores) that is used in several prior works would increase leakage power, and hence temperature, thus resulting in higher circuit-aging. In addition, higher leakage power dissipation would cause increased supply currents to be drawn from the PDN resulting in higher PDN-aging. In contrast, *ARTEMIS* prioritizes older (slower) cores that can support application frequency constraints without requiring hiking of V_{dd}-levels. The next section presents our problem formulation, followed by details of the *ARTEMIS* framework in Section 5.

4 Problem Formulation

4.1 Modeling BTI-Induced Circuit Aging

In this work, we model circuit aging effects arising from BTI-induced circuit degradation, as BTI has been found to be one of the most dominant aging mechanisms in emerging semiconductor technologies. But, our framework is capable of supporting models of other aging-mechanisms (HCI, TDDB etc.) as well. Velamala et al. [17], [18] have shown that a Trapping/Detrapping (TD) based BTI model is capable of accurately predicting the degradation under a sequence of V_{dd}'s used in the DVS operation. They have noted that when the supply voltage is changed from a higher V_{dd} to lower V_{dd}, the circuit-degradation undergoes recovery; this recovery behavior is not captured by conventional Reaction-Diffusion (RD) models. Therefore, our analysis of circuit-aging over the CMP-lifetime is based on the long-term aging prediction model proposed in [17], which accounts for different V_{dd}-levels over time. We estimate the effective ΔV_T increase that a component (computation core or NoC router) experiences over a time-interval of t using Eqs. (2) and (3):

$$\Delta V_T\ (t) = \ L.[A + B\log(1 + Ct)]. \tag{2}$$

$$L\ = K_1\ .\exp\left(\frac{-E_0}{kT}\right).\left\{exp\left(\frac{\beta V_1}{T_{ox}kT}\right).\alpha_1 + \cdots + exp\left(\frac{\beta V_S}{T_{ox}kT}\right).\alpha_S\right\}. \tag{3}$$

where V_i is the ith V_{dd}-level utilized by the component for a time-duration α_i, S is the total number of allowable V_{dd}-levels, and $\alpha_1 + \alpha_2 + \cdots + \alpha_S \leq t$. T is the average temperature of the component during corresponding α_i. We obtain A, B, and C and other parameter-values in Eqs. (2) and (3) by solving the equations given [1], [17], [18], [19] that are validated against silicon data. Values of constants used in Eqs. (2) and (3) are listed in Table 1.

4.2 Modeling EM-Induced PDN Aging

To model the phenomena of void nucleation and void growth in every horizontal and vertical on-chip wire, particularly those in the PDN that are under the most stress, we use the EM model proposed in [20] for copper (Cu) interconnects in the power grid. Eq. (4) gives the time t_n at which the void nucleates:

$$t_n = \frac{K_t}{D_{eff}},\quad K_t = \frac{\pi}{4}\ \left(\frac{(\sigma_c)^2\Omega k_B T}{\left(eZ_{eff}\rho j\right)^2 B}\right) \tag{4}$$

TABLE 1
Values of Constants Used in the Models of BTI and EM Aging

Constant	Description	Value
μ	Mobility of the charge carriers	1000 cm^2/V.s
α	Technology dependent constant	2
K_1	Poisson parameter for trap distribution	0.0075
E_0	Electric field across the channel	0.1897 eV
k, k_B	Boltzmann constant	0.000086
T_{ox}	Oxide thickness	1.4 nm
D_{eff}	Effective diffusivity	6.7×10^{-15}
σ_c	Critical stress	4.1×10^6 Pa
Ω	Atomic volume	1.182×10^{-29} m^3
ρ	Resistivity of Cu	2.5×10^{-8} Ω.m
B	Effective bulk modulus for the Cu-dielectric system	10^9
e	Charge of an electron	1.602×10^{-19} C
Z_{eff}	Apparent effective charge number	5.0
L_{wire}	Length of the wire	50 um

Once the void nucleates at time t_n, then at an observation time t_0, the length of the void L_{void} is:

$$L_{void}\ (t_0) = \left(\frac{D_{eff}}{k_B T}\right)\ eZ_{eff}\rho j(t_0 - t_n) \tag{5}$$

The length of the void in a Cu wire increases with the product of electrical current and the time-duration for which the current flows through it. The length of the void in turn determines the increased resistance (ΔR) or degradation of the wire, which is given by:

$$\Delta R\ =\ c.R_0\left(\frac{L_{void}}{L_{wire}}\right) \tag{6}$$

where R_0 is the resistance, constant c depends on the resistivity and cross-sectional area, and L_{wire} is the wire-segment length. Table 1 lists the values we used for the constants in Eqs. (4), (5), and (6).

4.3 Inputs, Assumptions, and Problem Objective

We have the following inputs to our problem:

- → A 3D NoC-based CMP with a 3D mesh NoC, of dimensions (dim_x, dim_y, dim_z) and number of tiles $N = dim_x \times dim_y \times dim_z$ with each tile containing a compute core and a NoC router;
- → A set S of candidate supply voltage (V_{dd}) levels for the chip;
- → A chip-wide dark-silicon power budget (DS-PB);
- → An application task graph for each application: vertices with task execution-times on compute cores and edges with inter-task communication volumes; execution time and volume values are assumed available from offline profiling;
- → Degree of parallelism (DoP) of each application, and a set of n admissible rectangular/cuboidal shapes (x, y, and z dimensions) of regions that it could be mapped to $B_1, \ldots, B_n$; e.g., a tuple set $2 \times 4 \times 1, 4 \times 2 \times 1, 2 \times 2 \times 2$ for an application with $\text{DoP} = 8$;
- → A minimum operating frequency (and thus a corresponding maximum execution time) constraint for each application;

→ A regular 3D power grid, with $p \times p$ grid-points supplying to each core and $dim_x \times dim_y \times p \times p$ power-pins at the top of the 3D-die, and an air-cooled heat sink at the bottom of the 3D-die;

We make the following assumptions in our work:

→ Applications can arrive in any order at runtime and must be mapped onto the 3D CMP while satisfying application-specific maximum execution time and minimum frequency constraints;

→ There exists one-to-one mapping between tasks and cores, i.e., a core can execute only one task (thread) at any given time;

→ Applications are mapped contiguously on non-overlapping rectangular (on single tier) or cuboidal (across multiple tiers) shaped regions of the 3D CMP for inter-application isolation and more optimized communication-profiles (as recommended in prior works such as [37]); thread-migration is not considered;

→ A chip-wide supply voltage exists that can be scaled using DVS at runtime, as the overheads of implementing DVS at a per-core granularity are high for CMPs with very large core counts [24];

→ Similar to prior-works (e.g., [3], [14], [16]) we assume presence of on-chip aging-sensors [22], [23] that provide aging information of compute-cores and NoC routers to our framework; we also assume voltage-sensors [21] at each power-input (PDN-grid-point) of a CMP-tile that track the severity of IR-drops (i.e., the degree of PDN-degradation on PDN-paths supplying to that core);

→ From the on-chip sensors, runtime sensed data (at a per-tile granularity), in terms of threshold-voltage (V_T) distribution and maximum IR-drops, is available after every *epoch*; our aging models (discussed in Sections 4.1 and 4.2) emulate runtime readings from sensors on real chips. We define an *epoch* as the time-period during which the aging profile of the chip can be assumed to be constant;

→ The 3D CMP is rendered unusable (end of lifetime) when an incoming application is unable to be executed (i.e., when its minimum application frequency requirement cannot be supported) on any of the allowed rectangular/cuboidal regions, at any V_{dd} level without violating the DS-PB constraint, when there are no other applications running at that time.

Objective: Given the above inputs and assumptions, our objective with the *ARTEMIS* framework is to perform run-time application-mapping and DVS-scheduling on a given 3D NoC-based CMP platform such that the total number of applications executed over the lifetime of the chip is maximized, while all application-specific minimum operating frequency- and maximum runtime-constraints, as well as CMP platform-specific DS-PB constraints are satisfied.

5 ARTEMIS FRAMEWORK: OVERVIEW

This section explains the design flow of the framework that involves the actual mechanism of run-time mapping and DVS selection along with the simulation of circuit- and PDN-aging in cores and routers. In manufactured CMPs, compact aging sensors such as the ones proposed in [22], [23], and [46] can be used to measure the on-chip aging profile. These sensors are amenable to use in standard cell design with minimal area and power overhead. They can be implemented in large numbers along the top 10 percent of component's critical paths to collect high volume digital data on device degradation.

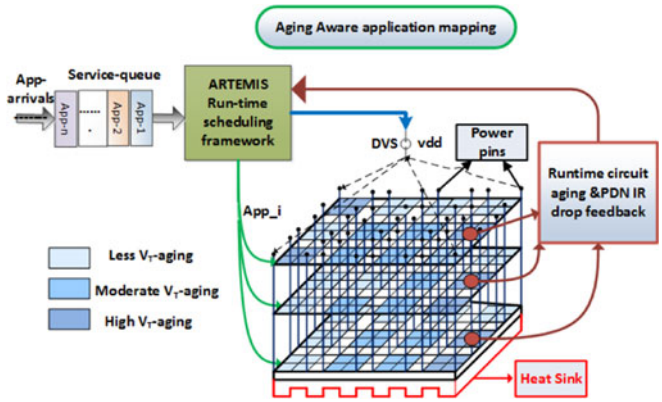

Fig. 2. Overview of *ARTEMIS* runtime aging-aware application-mapping and DVS-scheduling framework.

ARTEMIS is a run-time application mapping framework that processes applications from the top of the service queue, where the applications are stored as they arrive in the run-time. The run-time framework receives periodic feedback (at the end of each epoch) from on-chip aging and power sensors (simulated models), that is used for aging- and dark-silicon-aware mapping and DVS scheduling. At any given time, the scaling-down of V_{dd} via DVS-scheduling to save power and to limit aging is limited by the frequency-constraints of the applications running on the 3D NoC-based CMP, whereas scaling-up of V_{dd} is constrained by the DS-PB.

The key aspects of our proposed framework are illustrated in Fig. 2. The Run-time application-mapping consists of assigning the application task-graph on to a chosen rectangular- or cuboidal-shaped region of tiles on the 3D CMP admissible for the application (from the list $B_1, \ldots, B_n$), as well as performing routing path allocation of the intra-application communication-flows on the 3D NoC. The knowledge of the chip-aging profile is continuously utilized in the application-mapping and DVS scheduling steps.

The *ARTEMIS* framework is executed in two nested procedures: (i) Aging-aware application mapping and DVS (*inner-loop*); and (ii) Circuit- and PDN-aging analyses (*outer-loop*). These procedures are discussed in Sections 5.1 and 5.2 respectively, and the design flows for the two procedures are shown in Figs. 3a and 3b.

5.1 Aging-Aware Application Mapping and DVS Scheduling

During each epoch at runtime, we assume that applications arrive for execution on the 3D NoC-based CMP. Suppose a sequence of l applications arrives in an epoch. Aging-aware application mapping and DVS module (*inner-loop)* is responsible for mapping these l applications onto the CMP during the epoch. If an application cannot be mapped immediately after it arrives, it is kept in a service queue,

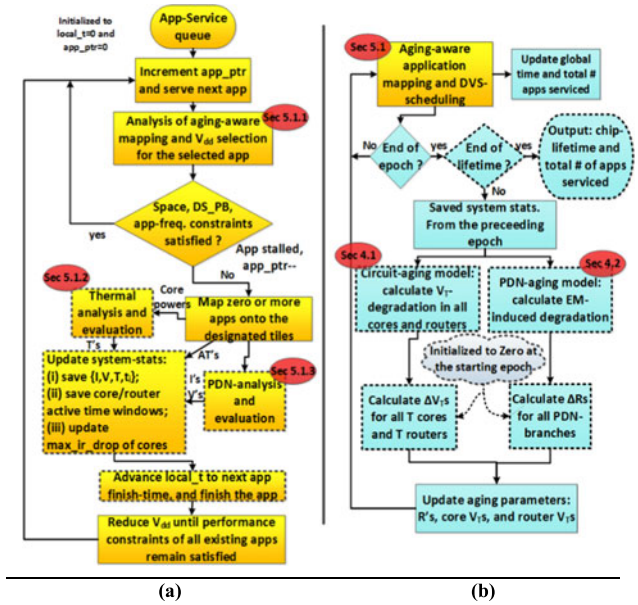

Fig. 3. *ARTEMIS* design-flow: (a) Aging-aware application-mapping and DVS scheduling (inner loop; Section 5.1). (b) Circuit- and PDN-aging analyses (outer loop; Section 5.2). The boxes with dotted outlines are used as part of our aging-simulation framework; however, these steps are not required on real hardware where runtime aging information is assumed to be available from on-chip sensors.

and mapped later. We assume processing of the service-queue on a first-come-first serve basis (although a priority-based processing approach could also be used). At the end of the epoch, aging information is updated from the on-chip circuit-aging sensors as well as the voltage sensors (this information is utilized by *inner-loop* during the next epoch). Subsequently, a new application sequence is serviced during each new epoch, and this process continues until the end of the lifetime of the 3D NoC-based CMP.

At the start of an epoch, the application-service queue is initialized to point to the first application ($app\_ptr = 0$), and the local time counter[1] is initialized to zero, as shown in Fig. 3a. Processing of the service-queue event is triggered, (i.e., new applications are serviced) when an application arrives or an existing one terminates. Once the event is triggered, an application instance is removed from the front of the queue and processed by the aging-aware mapping and V_{dd} selection phase (Fig. 3a; discussed in Section 5.1.1). Applications from the queue continue to be processed one-by-one until an "application stall" event is detected. An application in a service queue can be stalled only due to the following reasons: (i) available tile constraints on the 3D-die; (ii) DS-PB constraint; or (iii) application frequency constraints for the given degradation profile of the 3D-CMP. Note that if an application is stalled when there are no other applications running, i.e., the chip-degradation (V_T- and PDN-degradation combined) precludes it from meeting the application-frequency constraints, the 3D NoC-based CMP is considered as no longer usable and has reached its end of life.

When an "application stall" is detected or the application-service-queue becomes empty for the current epoch, the application(s) that have been processed by the mapping/selection phase (discussed in Section 5.1.1) are mapped on to the appropriate tiles chosen by the phase. At this time (local time), either one or more new applications are mapped on to the 3D-CMP or an application just ended (which triggered the service-queue), thus the steady-state computation-profile (i.e., CMP tile-power values, the resulting supply currents in the PDN, and thermal-profile) of the 3D-CMP changes. To evaluate the new computation-profile, given the tile-power-distribution, thermal-analysis is performed to re-evaluate the thermal-profile (at per-tile granularity) and PDN-analysis is performed to evaluate all the branch currents and voltage-drops at all grid-points in the PDN. Also, a worst-case IR-drop value (WC-IR-drop, which is the maximum voltage-drop out of all grid-points supplying to a tile) is evaluated for each of the N tiles. The WC-IR-drop value is updated for each tile (at every change of computation-profile) over the chip-lifetime and continuously used to calculate the maximum-frequency of the tile (for a given V_{dd}) in the application-mapping step. The thermal- and PDN-analysis is discussed in Sections 5.1.2 and 5.1.3, respectively. After the mapping/selection phase, thermal-analysis, voltage (V), and temperature (T) values, as well as the WC-IR-drop values in the time-window t_i for this (ith) computation-profile, is saved in the *system-stats*, as shown in Fig. 3a.

Additionally, if one or more applications are mapped at the current local time, the active-times (AT's) of compute-cores and NoC routers are calculated for each newly mapped application. For each tile, these AT's could be represented as $\{C_j, R_j, t_j\}$, where C_j and R_j take values of '1' or '0' depending on whether the corresponding compute-core or NoC router is active during the time-window t_j. These AT's for compute-cores and NoC routers are also saved in *system-stats*. The *system-stats* for all time-windows over the entire epoch duration are eventually utilized for aging-analyses (in the outer loop) at the end of the current epoch.

After updating *system-stats*, local time is advanced to the next application finish-time, and the corresponding application is completed, (Fig. 3a). As part of our DVS strategy to save power and limit aging, on completion of any application, we reduce V_{dd} to the lowest allowable level that would not introduce any violations in frequency constraints of existing (already running) applications.

5.1.1 Application-Specific Mapping and V_{dd}-Selection

For the application under consideration, this phase consists of three steps: (i) circuit- and PDN-aging aware region selection and V_{dd}-selection, (ii) communication-aware task-to-tile mapping, and (iii) NoC routing path allocation. We describe these steps below.

(i) Circuit- and PDN-aging aware region selection and voltage-selection: In our framework, an application with a given DoP can be mapped on to rectangular or cuboidal regions on the 3D CMP, with shapes to be chosen from a pre-defined list $\{B_1, \ldots, B_n\}$ for that application. All intra-application communication is contained within t closed region, thus application-isolation is maintained and communication cross-interference

1. For simulation purposes, we keep track of cumulative execution-times of cores and routers of the CMP, by utilizing a local time counter for the *inner loop* that is reset at the start of each epoch, and a global time *outer loop* that is augmented by the local time at the end of each epoch. In an actual system, however, the local time counter is used to keep track of the epoch durations. At the end of each epoch, the aging analysis module is invoked that computes the new circuit- and PDN-aging profile of CMP.

is eliminated. Our heuristic in this step utilizes the V_T-degradation profile and the WC-IR-drop profile of the 3D CMP. The objective is to find the region on the 3D-mesh (with one of the admissible shapes) so as to: *(a)* minimize leakage-power; *(b)* minimize EM-induced degradation of PDN-paths supplying to cores with high WC-IR-drops; *(c)* satisfy the frequency-constraint of the application by all cores within the region; *(d)* satisfy the DS-PB. In other words, we search for CMP-regions with most circuit-aging that satisfy minimum application-frequency constraints and have least WC-IR-drops. To this end, we define the following cost-function (Ψ) for joint optimization of leakage-power and PDN-degradation:

$$\Psi = \sum_{k=1}^{k=DoP} \left\{ \alpha . \left(\frac{\max\_V_T - V_{Tk}}{\max\_V_T - nom\_V_T} \right) + \beta \left(\frac{WC\_IR\_drop_k}{\max\_IR\_drop} \right) \right\} \tag{7}$$

where, V_{Tk} is the effective V_T and $WC\text{-}IR\text{-}drop_k$ is the WC-IR-drop of the kth core within the region of DoP cores; nom_V_T is the nominal (lowest) effective V_T-value of a core with no aging; and α and β are weighting coefficients. We define max_V_T as the maximum V_T value that the core can support for an ideal (zero) WC-IR-drop (at highest V_{dd}) while meeting the frequency-constraint of the application. Similarly, *max_IR_drop* is the maximum tolerable WC-IR-drop for a core for nominal V_T and highest V_{dd}.

Algorithm 1. Aging-aware Region Selection and V_{dd}-Selection Heuristic

Inputs: ***V_T-profile, WC-IR-drop profile***, $\{B_1, \ldots, B_n\}$
while ($V_{dd} \leq max\_V_{dd}$) **do** {
  for each tile on the 3D NoC-based CMP **do** {
    assume this tile to be at the minimum x, y, z coordinates of the region
    for each shape in $\{B_1, \ldots, B_n\}$ **do** {
      if all tiles (compute-cores and routers) satisfy app-frequency
      check if DS-PB is satisfied
        calculate Ψ, choose this shape if least Ψ **AND** DS-PB satisfied
      else go to next shape B_i (step 4)
      end if
  } // end for each shape ...
} // end for each tile ...
if (no valid region found **AND** no DS-PB violation)
  hike V_{dd}
end if
} //end while
if no valid region found
  stall this application
end if
output*: *a valid region to map the application and V_{dd}-level, or "stall"

Algorithm 1 shows the pseudo code of our region- and V_{dd}- selection heuristic. The heuristic performs a simple exhaustive search over all tiles on the 3D-mesh and over all admissible shapes $B_1, \ldots, B_n$ for the application under consideration. The V_T-profile and WC-IR-drop profile inputs are used for calculating the value of Ψ. The region with the least Ψ value that satisfies the frequency-constraints (with maximum frequency for the selected V_{dd} level calculated using Eq. (1)) and at the same time does not violate the DS-PB (given that existing applications have been running), is selected for mapping the application under consideration. If no region on the 3D-mesh is found to satisfy the frequency-constraints, we repeat the search for successively higher V_{dd}-levels (which can allow using a higher frequency as per Eq. (1) with a better probability of meeting frequency-constraints) until either a valid region with minimal Ψ is found or the DS-PB is violated. If no valid region is found, an "application stall" event is initiated.

We now present the theoretical time-complexity of this heuristic. At most N tiles (total tiles on the 3D NoC-based CMP) are considered for the prospective mapping region. Note that DoP of the application (relatively small integer c – treated as a constant) number of tiles are to be evaluated for frequency and leakage-power at each of these iterations. As, the number of candidate V_{dd}-levels $|S|$ as well as the number of admissible shapes n are expected to be small constant integers, our region-selection step runs in linear complexity with respect to the number of tiles, N: O($cn|S|N$).

(ii) Communication-aware task-to-tile mapping: After the region on the 3D CMP has been selected (of size equal to application-DoP), our mapping heuristic maps the appropriate application-task-graph on to the chosen CMP tiles. We utilize a fast and efficient communication-aware incremental-mapping approach (similar to that used in prior works such as [26], [27]) suitable for runtime use.

(iii) Symmetric aging-enabled routing path allocation (SAR): In this step, we map the communication-flows of the current application on to the designated cuboidal region on the 3D NoC-based CMP. We propose an aging-enabled and congestion-aware routing scheme (SAR) to produce a balanced core-router aging profile and extend the lifetime of the 3D NoC. The main objective of SAR is to minimize the number of runtime scenarios where application-mapping on a given cuboidal region is precluded due to aging in routers. Note that an application can be mapped only if all tiles (each tile has a compute-core and a NoC-router) within the region under consideration satisfy the minimum application-frequency constraint. Prior work on aging- enabled routing (such as [35]) considers the aging in NoC-routers but does not consider the aging in compute-cores. Such an approach could lead to a somewhat imbalanced aging within tiles of the CMP, thus potentially preventing application mapping onto desirable CMP regions due to excessive aging in NoC routers. SAR on the other hand enables symmetric aging on individual tiles of the 3D-CMP to extend the service life of NoC routers. Additionally, SAR efficiently trades-off aging with network-congestion in the NoC by selecting routing paths to maximize NoC-lifetime while leveraging the knowledge of maximum execution time constraints of applications, i.e., the aging metric in the routing cost function is prioritized by varying degrees, given the time-slack available for application-completion.

To ensure a low-overhead implementation, path diversity, and deadlock freedom, our routing algorithm builds on the 4N-First turn model [36] for 3D-mesh NoCs. This routing algorithm is partially adaptive, and hence allows the flexibility to potentially select from among multiple next hop directions, at each router. We designed a cost-function

for next-hop selection during routing that considers the difference between router-aging and core-aging (*router_V_T – core_V_T*) values to ensure balanced aging in CMP tiles. Moreover, as congestion in the NoC-links leads to excessive routing delays and thus longer application-runtimes, we prefer allocating flows to links with lesser communication-volumes. The following routing cost function (Rt_{cost}), which is a linear combination of the two normalized metrics, is used to make routing decisions at each hop along the path:

$$Rt_{cost} = \alpha_R.\frac{(V_r\ difference) - (minimum\ V_r\ difference)}{range\ of\ V_r\ difference} + \beta_R.\frac{(volume) - (minimum\ volume)}{range\ of\ volume}, \quad (8)$$

where, α_R and β_R are weighting coefficients, V_T *difference* represents (*router_V_T – core_V_T*) of the candidate next hop router, and volume represents the existing communication-volume (already allocated while routing previous flows) on the link. SAR selects the next hop with the minimum routing-cost, Rt_{cost}, given in Eq. (8).

Algorithm 2. Symmetric Aging-Enabled Routing Path Allocation

Inputs: Task-graph, execution time constraints, minimum frequency, task-mapping of current application, V_T–profile of compute-cores and routers

1: **Initialize** $\alpha_R = 1$ and $\beta_R = 0$
2: **for** all communication-flows **do** {
3: **for** all hops on the minimal path **do** {
4: select the next hop with the least Rt_{cost} (Eq. (8))
5: } update α_R and β_R (Eq (9))
6: }

output: all flows of the application allocated on the cuboidal CMP- region

Communication delays are calculated from the application-frequency and NoC link bandwidths, and thus the current application-delay can be estimated from the already routed communication-flows. NoC routers and links in an application region run at the same frequency as the cores in the region (application-frequency). Note that the goal of SAR is to extend NoC lifetime while meeting application execution time constraints. Thus, the values of coefficients in Eq. (8), α_R and β_R, are re-evaluated after routing each flow, as shown below:

$$\beta_R = \{current\ app.\ delay\}/\{\delta.(app.\ execution\ time\ constraint)\}$$
$$\alpha_R = (1 - \beta_R). \quad (9)$$

Before any application communication flows are mapped to the NoC routers, we start with values $\alpha_R = 1$ and $\beta_R = 0$. As flows are mapped and the estimated application-delay increases, the value of β_R increases (α_R decreases) proportionally until the application-delay reaches a significant fraction (δ) of the application execution time constraint. At this point ($\beta_R = 1$ and $\alpha_R = 0$), SAR ceases to be aging-aware and routes on paths with minimum congestion exclusively, to meet the execution time constraint of the given application. Algorithm 2 below summarizes our symmetric-aging enabled routing scheme.

Note that the given application is executed on the actual 3D-CMP platform only after the analysis for routing path allocation is performed. The turn model rules are implemented in each router using simple combinational logic. The next hop selection information at each NoC router is stored in small next-hop routing tables that enable quick selection of the most appropriate next-hop direction based on the source and destination of a packet. Even for the largest sized, 32-threaded applications mapped onto a {4x4x2} cuboid on the 3D CMP platforms we considered, we found that the upper bound on number of communication-flows (with unique source-destination pairs) needed to be routed through any router is 64, with our 3D turn-model based minimal routing scheme. Thus, a NoC router on the 3D CMP would need a next hop table of up to 64 entries. Assuming 3 bits for the output port and 6 bits for the source and destination each, the footprint of the NoC routing table is only 960 bits. As we consider communication intensive applications for execution on the 3D CMP, there is a need for deeper buffers at input and output channels to avoid severe network congestion and application slowdown. For such conditions, we provide each input/output channel with four virtual channels, each consisting of a buffer of size four flits. Hence, the overall size of the buffers is up to 1.7KB, assuming the flit size to be 8B. Thus, the hardware overhead of implementing SAR is small (960 bits or 0.12KB) when compared to the total size of the buffers.

5.1.2 Thermal-Analysis and Evaluation

To perform thermal evaluation of a given computation-profile in our framework, we utilize the open-source thermal emulator 3D-ICE 2.2.5 [29] which supports steady-state thermal analysis of 3D ICs with a conventional air-cooled heat-sink. For the given power-profile, the tool outputs the core-temperatures (T's) on the 3D die.

5.1.3 PDN-Analysis and Evaluation

The supply current drawn by each core is calculated from the core-power and selected V_{dd}-level. Given the supply current requirements of the N cores on the 3D-CMP, we created a linear programming (LP) formulation and used lp_solve [28] to solve for the grid-point voltages and currents flowing in the 3D regular power grid. This enables the updating of {V's, I's} in the power-grid and WC-IR-drops of cores in the 3D CMP, for the given time-window (t_i) of the computation-profile. A more detailed discussion of our LP formulation, including elaboration of constraints and equations, is presented in Appendix I, which can be found on the Computer Society Digital Library at http://doi.ieeecomputersociety.org/10.1109/TMSCS.2017.2686856.

5.2 Circuit- and PDN-Aging Analyses

In the outer loop of our framework (Fig. 3b), we utilize *system-stats* generated by the inner loop over the last epoch to perform aging-analysis at the end of the epoch. Given the *system-stats* for the last epoch, this analysis is used to calculate the rise in effective V_T values (ΔV_T's) of all cores and NoC routers on the 3D CMP, as well as the rise in resistance values (ΔR's) of all vertical and horizontal PDN-branches, using the circuit- and PDN-aging information. The BTI-

induced circuit-aging of compute-core and NoC router components are calculated (discussed in Section 4.1) using the V's and T's experienced by these components during all of their AT's over an entire epoch. The EM-induced PDN-degradation in PDN-branches (discussed in Section 4.2) is calculated using I's for all computation-profiles of the epoch. As effects of EM are far less dominant in signal interconnects compared to PDN-interconnects [4], [12], we ignore EM-induced aging in the NoC-links and focus primarily on PDN interconnects.

Note that the active-time windows of compute-cores and routers $\{C_j$'s, R_j's, t_j's$\}$ may not be aligned with the chip-wide computation-profile windows $\{V$'s, T's, I's, t_i's$\}$; therefore, in circuit-aging calculations, the component AT's are required to be split into multiple time-windows where computation-profiles change. Also, at the start of the very first epoch, the R's and V_T's are initialized with nominal values representing no degradation and the ΔR's and ΔV_T's are initialized to zero-values. Lastly, the updated aging profiles are leveraged to make mapping decisions in the next epoch. When the end of lifetime is encountered (discussed in Section 5.1), the aging analyses procedure outputs the lifetime of the 3D-CMP in terms of both the total system-execution-time (global time) and the total number of applications serviced during this time (Fig. 3b).

6 Experimental Studies

6.1 Experimental Setup

Our experiments were conducted using 13 different parallel application benchmarks taken from the well-known SPLASH-2 [31] and PARSEC [32] benchmark suites. We profiled the execution-time, power dissipation, and degree of memory-intensity of each application for different application-DoPs by performing multicore simulations using the open-source tools SNIPER [33] and McPAT [34]. For each benchmark, the DoP resulting in highest performance was obtained from this profiling study and selected as the fixed DoP value for that benchmark. These DoP values ranged from 4 to 32. Note that increasing DoP beyond this baseline value for each benchmark resulted in lower performance, due to inter-thread synchronization and communication overheads.

We categorized the 13 benchmarks into two groups: (i) communication-intensive benchmarks - {*cholesky, fft, radix, raytrace, dedup, canneal, and vips*}; and (ii) compute-intensive benchmarks—{*swaptions, fluidanimate, streamcluster, blackscholes, radix, bodytrack, and radiosity*}. As *radix* has properties of both, we use it in both groups. In our analyses, we employ three types of application sequence groups as inputs to our framework: communication-intensive, compute-intensive, and mixed (using all 13 applications). We assume each application-sequence to have 100 randomly ordered application-instances selected from the respective group. To enhance the statistical significance of our results, we averaged results for five different randomly generated application-sequences for each group.

To simulate the chip-lifetime within a reasonable time, we extrapolate the effects of aging over 500 such sequences, making the total number of application-instances executed within an epoch to be approximately $l = 50{,}000$. Simulation times for *ARTEMIS* to simulate till the end of the lifetime were between 6 and 10 hours. The communication-intensive application workloads typically entailed larger simulation times because of longer chip lifetimes (see results and discussion in Section 6.2), compared to the computation-intensive workloads.

We consider a 60-core 3D-mesh NoC based CMP platform, with dimensions $5{\times}4{\times}3$ ($dim_x{\times}dim_y{\times}dim_z$). Our SNIPER simulations for application-profiling capture performance and power consumption at the 22nm process technology node. Seven operating voltage levels are used, ($|S| = 7$): 0.7 V, 0.75 V, 0.8 V, 0.85 V, 0.9 V, 0.95 V, and 1.0 V. Frequency-requirements of different applications are set between 1.5 GHz and 2 GHz. The following region-dimensions-lists $\{B_1, \ldots, B_n\}$ for applications (for the given DoPs) are employed: $\{2{\times}2{\times}1\}$ for $\mathrm{DoP} = 4$, $\{4{\times}2{\times}1, 2{\times}4{\times}1, 2{\times}2{\times}2\}$ for $\mathrm{DoP} = 8$, $\{4{\times}2{\times}2,\ 2{\times}4{\times}2,\ 4{\times}4{\times}1\}$ for $\mathrm{DoP} = 16$, and $\{4{\times}4{\times}2\}$ for $\mathrm{DoP} = 32$. The dark-silicon power-budget (DS-PB) is conservatively set at 85W. The regular 3D-PDN power grid is modeled based on guidelines provided in [30]. With 20 cores on each tier, a total of 320 input power pins are used with $n^2 = 16$ grid-points for each core. Nominal (initial non-aged) values of branch resistances are assumed to be 50 mΩ [30], with 25 μm^2 cross-sectional area.

For our circuit-aging calculations, we assume a nominal effective V_T of 0.3 V for un-aged cores and routers. In our combined cost function calculations (Ψ in Eq. (7)) for the aging-aware region-selection heuristic, we use $\alpha = \beta = 0.5$ (empirically derived to achieve the longest lifetimes); *max_IR_drop* and *max_*V_T are set to 0.3 V and 0.5 V respectively, based on Eq. (1), with operating frequency requirement of 2 GHz. In our SAR heuristic, we use $\delta = 0.6$, to calculate the value of β_R, for an appropriate trade-off between application performance and aging. In our experiments, an epoch interval can range between 25 to 35 days, depending on the power profile, execution-times, and average DoPs of the application workload, as well as the degree of aging in the chip. Given the relatively slow rate of aging, such an aging-measurement interval has been found to be appropriate for runtime frameworks [3]. Also, as the overheads incurred due to employing aging sensors have been reported to be quite small (power dissipation of 84.7 nW, sensing-latency of 100 μs, and area of 77.3 μm^2 per sensor at 45 nm technology node) in [22], we ignore them in our calculations.

6.2 Experimental Results

Our experiments compare three variants of the proposed *ARTEMIS* framework with two other runtime mapping approaches derived from prior work. These prior works are designed for 2D CMPs, so we extend them to 3D CMPs for a fair comparison.

To investigate the effectiveness of the circuit-aging (leakage) and PDN-aging aware region-selection, voltage-selection, and mapping techniques, we adapt our *ARTEMIS* framework to use an XYZ-routing scheme (*ARTEMIS-XYZ*) and compare the results obtained with two other run-time mapping techniques that selects contiguous regions for mapping and use the same XYZ routing scheme: (i) *traditional worst-case guard-banding approach (WC-GB)*: In this approach, region selection is done based on the runtime area constrained mapping approach from [37] that attempts to fit the maximum number of applications on the chip. To satisfy the

application-performance requirements for an extended period of time, a high $V_{dd} = 1.0$ V is used at all times. This framework selects contiguous regions for mapping, to maximize the performance, and minimize the communication latency. However, it does not assume runtime inputs from aging-sensors to make mapping decisions and thus is not aging-aware; (ii) *wear-leveling approach with DVS (WL+DVS)*: In this approach, contiguous region-selection for application-mapping is always done based on the lowest average V_T-degradation in cores, as proposed in [13], [14]; in addition, V_{dd} is opportunistically reduced when possible and adaptively hiked with aging to meet application performance constraints.

Additionally, we adapt our *ARTEMIS* framework to use an aging- and congestion-aware routing scheme (ACR) obtained from prior work in [35]. We also include results for our *ARTEMIS* framework with the proposed symmetric aging-enabled routing (SAR) scheme. Thus, the comparison between *ARTEMIS-XYZ*, *ARTEMIS-ACR*, and *ARTEMIS-SAR* allows us to determine the most effective 3D NoC routing approach that can help improve lifetime in 3D NoC-based CMPs while meeting application performance and chip-wide power constraints. Finally, to test the efficiency of the ARTEMIS framework, all the experiments are conducted when the workload is high at a uniform inter-application arrival rate of ~1.5s while each application executes for few seconds on the 3D CMP.

Fig. 4a shows the total number of applications serviced over the chip lifetime, Fig. 4b shows total CMP-lifetime (total system-execution-time), and Fig. 4c shows the application-throughput extracted over the service-life of the CMP for all the compared frameworks, across the three different types of application-input-sequences. The results shown in Fig. 4 are the error-bars in all our plotted results represent the range of results across simulations with five different randomly generated application-sequences (with individual applications in the sequence derived from the SPLASH-2 and PARSEC benchmark suites, as discussed earlier).

As expected, the WL+DVS framework outperforms the WC-GB approach that does not perform DVS. By intelligently selecting application-regions on the 3D-die with its region-selection heuristic, our *ARTEMIS* frameworks (*ARTEMIS-XYZ*, *ARTEMIS-ACR*, and *ARTEMIS-SAR*) achieve a notable reduction in leakage-power dissipation and reduce stress on the more highly degraded PDN-paths. The *ARTEMIS* frameworks produce 9–40 percent (25 percent average) improvement in the total number of applications serviced over the next best framework, WL+DVS, as well as significant improvements in total CMP-lifetime, as can be seen from Figs. 4a and 4b. For communication-intensive applications, we observed far less percentage of dark-silicon, approximately 0–15 percent (depending on V_{dd}-levels and workload profiles), compared to compute-intensive applications where dark-silicon is approximately 10–33 percent. A lower percentage of dark-silicon is indicative of more active cores running with less stress, whereas a higher percentage of dark-silicon indicates fewer active cores that are running with greater stress.

Thus, communication-intensive applications experience less aggressive aging (because of their lower %dark-silicon), which results in more of these applications being executed over the chip lifetime and a higher lifetime compared to compute intensive applications. Figs. 4a and 4b corroborate this observation. Also, most communication-intensive applications generate relatively low current-densities in the PDN, i.e., PDN-degradation is slower relative to circuit-degradation, which limits the improvements obtained by the *ARTEMIS* frameworks for such applications, as can be observed from Fig. 4a.

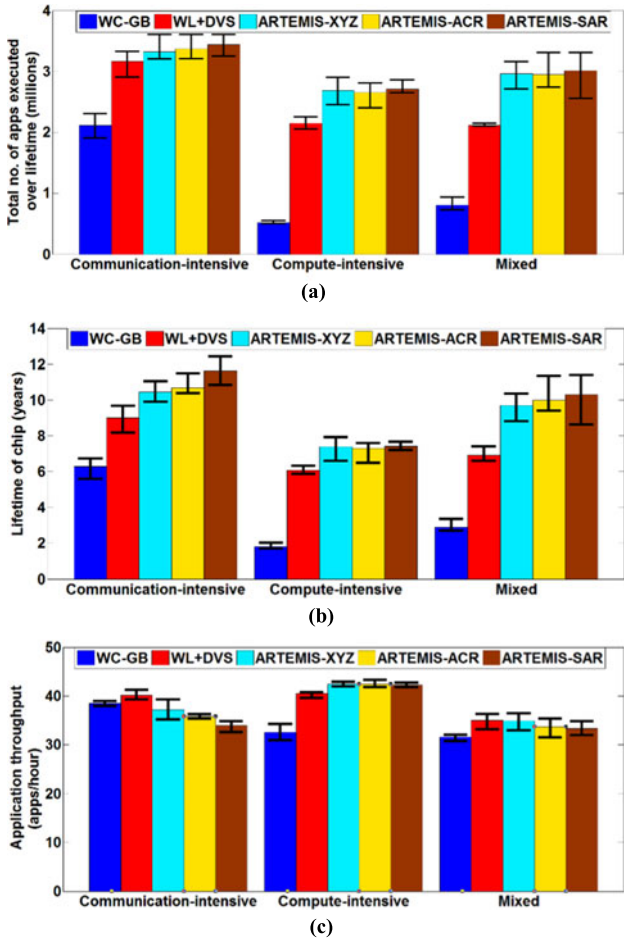

Fig. 4. Results comparing *ARTEMIS* framework variants with other approaches from prior work, for workloads that combine various SPLASH-2 and PARSEC benchmarks: (a) Total number of applications serviced over lifetime, (b) lifetime (years), and (c) application-throughput over lifetime (applications/hour).

Next, we present an analysis of lifetime improvements obtained when our proposed symmetric aging-aware routing path allocation (SAR) heuristic is used with *ARTEMIS* (*ARTEMIS-SAR*), compared to the *ARTEMIS-XYZ* and *ARTEMIS-ACR* frameworks. Our SAR heuristic enables better balancing of aging between compute-cores and their associated NoC-routers. While executing communication-intensive workloads exclusively, where the rate of aging in routers is comparable to that of core-aging, SAR minimizes the number of runtime scenarios when mapping of an application is stalled due to aged routers, thereby extending the system-lifetime. Observe in Fig. 4a that *ARTEMIS-SAR* produces notable improvements in number of applications executed over lifetime, compared to *ARTEMIS-XYZ* (by 4 percent) and *ARTEMIS-ACR* (by 2.2 percent) with comparable application-throughput (as shown in Fig. 4c), for communication-intensive workloads. However, the choice of

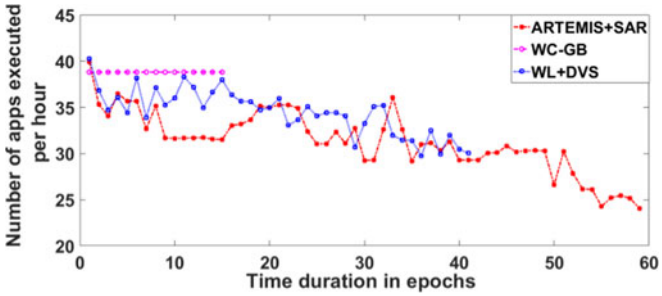

Fig. 5. Results showing the comparison of application throughput of *ARTEMIS+SAR*, WL+DVS, and WC-GB over their respective CMP lifetimes.

routing scheme has very little effect on lifetimes for compute-intensive and mixed workloads, where NoC-aging does not determine the service-life of the chip.

To further analyze how the behavior of different frameworks change over time, we show how application throughput varies across time in different frameworks in Fig. 5. We consider *ARTEMIS* with *SAR* routing scheme and a mixed workload in this experiment. It can be seen from Fig. 5 that WC-GB maintains a constant application throughput, while WL+DVS and *ARTEMIS+SAR* tradeoff throughput for an extended CMP lifetime. WL+DVS and *ARTEMIS+SAR* employs intelligent aging-aware application mapping for a graceful degradation of the CMP, and utilizes DVS to satisfy the application minimum frequency and DS-PB constraints. *ARTEMIS+SAR* further benefits from PDN-aging-aware intelligent mapping scheme that extends the useful lifetime of the CMP beyond that of WL+DVS, meanwhile achieving similar throughput.

We also show experimental results related to the power dissipation, PDN performance, and V_T degradation profile on the 3D CMP when using different mapping frameworks, in Fig. 6.

A comparison of the average power-dissipated per application over the chip lifetime is shown in Fig. 6(a). As expected, WC-GB framework, which does not utilize DVS, dissipates significantly more power. The leakage-optimizing mapping in *ARTEMIS* results in up to a 5.5 percent improvement for compute-intensive workloads (2.8 percent on average for all workloads) in total power/application over WL+DVS. We also analyze the distribution of percentage worst-case IR-drops (%WC-IR-drops) at the end of lifetime with different frameworks. Fig. 6b shows the maximum %WC-IR-drops obtained for different frameworks at the end of chip lifetime. The aging-unaware WC-GB framework maps applications such that some cores are more heavily loaded than others, thus resulting in the shortest lifetimes with high maximum WC-IR-drops. With our strategy to prioritize mapping on cores with less WC-IR-drops, *ARTEMIS* frameworks produce lower maximum %WC-IR-drop-values (by up to 9 percent lower), compared to WL+DVS, despite *ARTEMIS* having a longer lifetime and servicing a higher number of applications.

Fig. 6c shows the variance in the WC-IR-drop-distribution on the 3D chip obtained at the end of lifetime with different frameworks. A smaller variance of IR-drops with *ARTEMIS* frameworks (up to 24 percent lower compared to WL+DVS) signifies efficient management of PDN-aging that aides in improving the longevity of the PDN, and thus the entire chip.

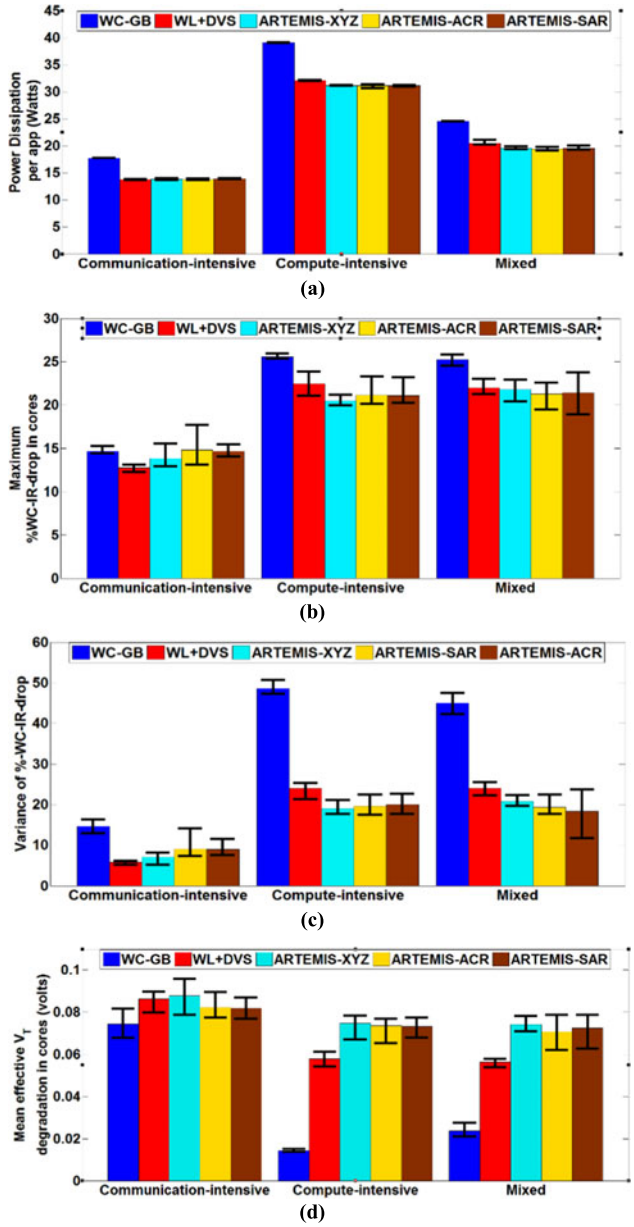

Fig. 6. Results showing improvements for our circuit-aging (leakage) and PDN-aging aware region selection and V_{dd} selection heuristic in the *ARTEMIS* frameworks: (a) power dissipation per application, (b) maximum percentWC-IR-drop at end of lifetime, (c) Variance of percent WC-IR-drop at end of lifetime, and (d) Mean effective V_T-degradation in compute-cores at end of lifetime.

Fig. 6d, shows the mean effective V_T-degradation in compute-cores at the end of lifetime, and provides additional insights into the lifetime improvements obtained with our circuit-aging (leakage) and PDN-aging aware region selection and V_{dd} selection heuristic. As discussed earlier, the V_T-values of circuit components increase with aging. Given the nominal-V_T of 0.3V at the start of lifetime, observe in Fig. 6d that the mean V_T-degradation values at the end of lifetime for *ARTEMIS* frameworks are significantly higher (by up to 30 percent for compute-intensive workloads) compared to the WL+DVS framework. By restricting the EM-induced PDN-degradation, *ARTEMIS* can extend the tolerable degree of circuit-aging (V_T-degradation) in compute-cores, while meeting the same performance constraints. Thus, the 3D CMP remains functional

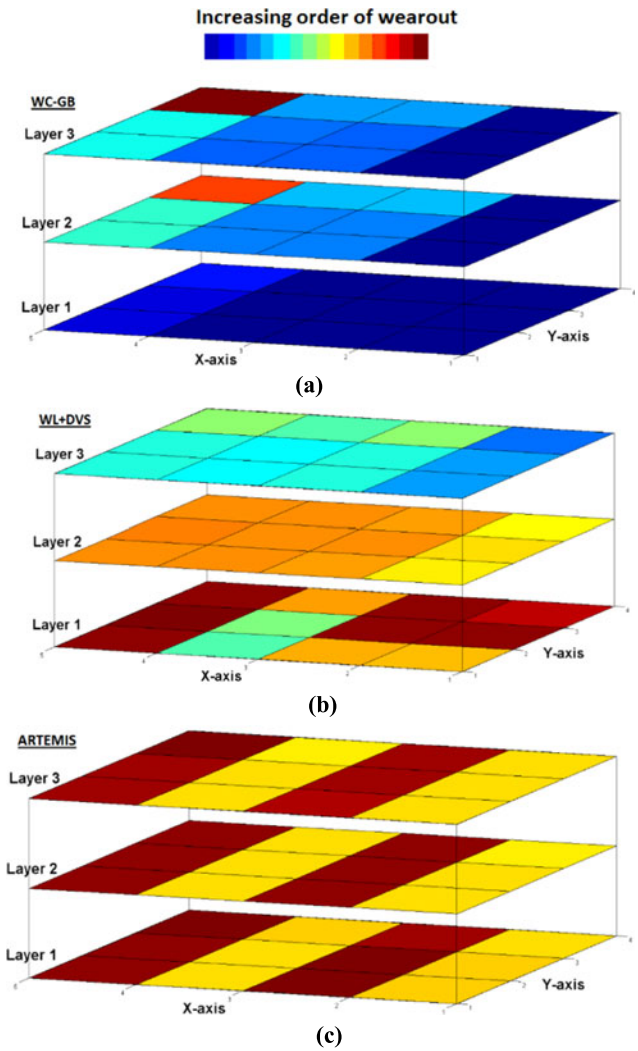

Fig. 7. Surface plots showing the effective V_T degradation in cores of a 3D CMP at the end of their respective lifetimes using: (a) worst case guard-banding (WC-GB) technique, (b) wear leveling with DVS (WL+DVS) technique, and (c) proposed ARTEMIS-SAR framework.

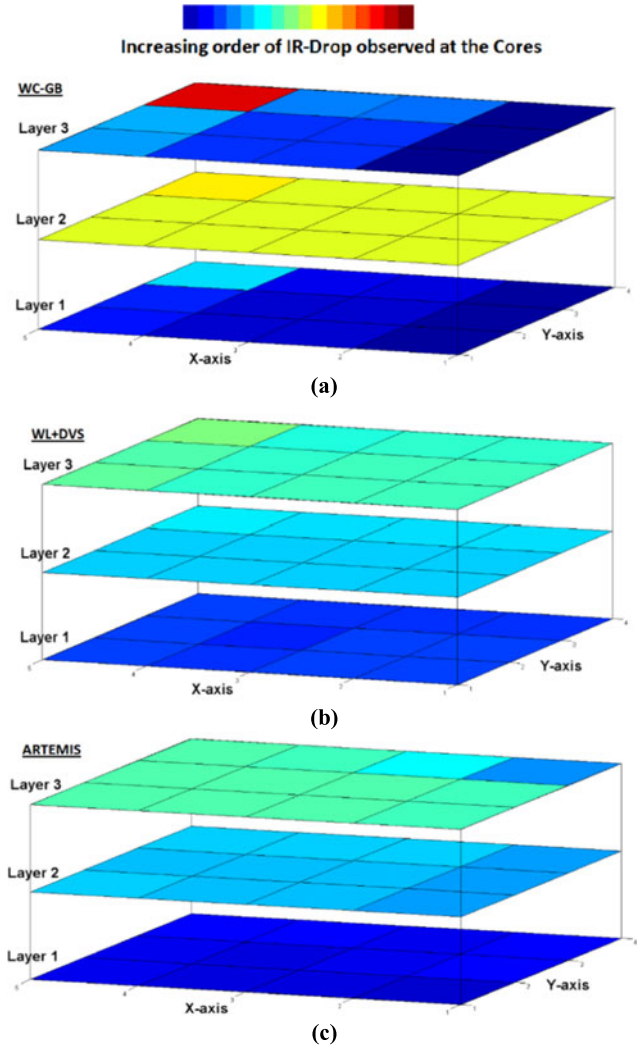

Fig. 8. Surface plots showing the Worst-Case IR drops observed on each layer of a 3D CMP at the end of their respective lifetimes using: (a) worst case guard-banding (WC-GB) technique, (b) wear leveling with DVS (WL+DVS) technique, and (c) proposed ARTEMIS framework.

for much higher V_T-degradation with *ARTEMIS* compared to other approaches.

To obtain a more comprehensive understanding of the performance of the compared frameworks, we present snapshots of the 3D-CMP die at the end of lifetime, when using different frameworks. Fig. 7 shows average values of V_T-degradation and Fig. 8 shows maximum WC-IR-drops in cores at the end of CMP lifetime after executing compute intensive workloads. We have only considered compute intensive workloads for this analysis because V_T-degradation and IR-drops due to these workloads are higher in cores due to their higher power dissipation (compared to mixed and communication-intensive workloads).

Figs. 7a, 7b, and 7c show the average V_T-degradation observed in each core at the end of CMP lifetime for WC-GB, WL+DVS, and *ARTEMIS-SAR* frameworks. We consider the end of lifetime as the condition when application performance constraints are not met at any CMP V_{dd} level, without violating the chip's dark-silicon power budget (DS-PB) constraint, and when no other application is currently running on the CMP. Fig. 7a shows that when WC-GB fails, most of its cores are comparatively less aged than that of WL+DVS in Fig. 7b and *ARTEMIS* in Fig. 7c. This is because WC-GB ambitiously tries to map the applications at a very high V_{dd} repeatedly on to the same, rapidly aging cores. Also, the PDN degradation in WC-GB is much higher than the other two frameworks in Fig. 8a. As the degradation in PDN worsens, application performance constraints are not met within the allowed set of CMP V_{dd} levels. This results in end of lifetime condition in WC-GB, because when application-performance constraints are not met in a region, WC-GB does not exhaustively search for other mapping regions on the CMP even if no other applications are running simultaneously. Fig. 7c shows that by using *ARTEMIS*, the chip has more degraded cores at the end of the lifetime compared to the other two frameworks. This indicates that all the cores have been well utilized by *ARTEMIS* till the time it failed. Better management of EM-degradation in PDNs (which is explained using Fig. 8), and trading off application throughput for lifetime by *ARTEMIS* (shown in Fig. 5) helps to extend the functional lifetime of each individual core on the 3D CMP till the end of its lifetime.

Also, note that with *ARTEMIS*, cores are degraded in a pattern that is beneficial for contiguous mapping of

TABLE 2
Primary Reason for Failures Across Different Frameworks

	WC-GB	WL+DVS	*ARTEMIS*
Primary reason for failure	Core + PDN aging	PDN aging	Core aging

applications till the CMP fails, unlike with WC-GB and with WL+DVS that leave fragmented regions of cores with varied degradation profiles.

Figs. 8a, 8b, and 8c show the Worst-Case-IR-drops (WC-IR-drops), due to EM of PDN, observed at the end of the CMP lifetime. The colors at each tile give a comparative visualization of IR-drop values logged at the input pins of the cores. From Figs. 8a, 8b, and 8c, it can be seen that WC-GB is completely unaware of PDN aging while mapping, and hence has some hotspots (red and brown tiles) due to excessive mapping of tasks on to the same tiles, even when other tiles are free and less degraded. WL+DVS show slightly higher WC-IR-drops to *ARTEMIS* at the end of the CMP lifetime. By integrating PDN-awareness and simultaneously managing V_{dd} and mapping regions intelligently, *ARTEMIS* makes it possible to use the CMP well beyond the V_T*-degradation* values (with similar PDN-degradation) observed using prior works (as shown in Figs. 7a, 7b, and 7c).

The aging aware mapping heuristic WL+DVS that is unaware of PDN degradation, tries to map applications on to tiles with less circuit degradation. Hence, the PDN ages at a faster rate in WL+DVS compared to *ARTEMIS*. But, as PDN degradation gets higher towards the end of the lifetime, applications do not meet their frequency constraint, and force a CMP V_{dd} hike at the time of mapping. Also, cores tend to dissipate more power as they get older, leading to DS-PB violation. This results in CMP reaching the end of lifetime condition faster. Hence PDN-aging is crucial for achieving longer lifetime in 3D CMPs. *ARTEMIS* is thus able to manage both circuit and PDN aging to extract more work out of the CMP in its lifetime.

Table 2 summarizes the primary reason for failure in WC-GB, WL+DVS and *ARTEMIS*, as observed in Figs. 7 and 8. WC-GB is aging unaware, leading to an exacerbated aging of cores and PDN. Hence, the primary reason for failure, and the bottleneck for achieving longer lifetime in WC-GB is both core and PDN aging. WL+DVS prioritizes mapping on the younger cores with no knowledge of PDN degradation, which makes PDN aging a bottleneck for achieving better lifetime, and a primary reason for failure. ARTEMIS balances core and PDN aging to extract the useful lifetime from the CMP. Hence, in ARTEMIS the end of lifetime condition is reached due to core aging.

7 Conclusion and Future Work

In this paper, we proposed an aging-aware application-mapping and DVS scheduling framework (*ARTEMIS*) that considers PDN-aging of 3D NoC-based CMPs in addition to circuit-aging (in NoC routers and cores) in both the performance and aging evaluation stages, and the lifetime optimization methodology. We have considered the analysis of *ARTEMIS* framework in a highly-constrained system with variable application execution time. Compared to a framework based on the best known prior work on aging-aware mapping techniques, *ARTEMIS* can service 25 percent more applications (on average) over the chip lifetime, which highlights its promise for emerging 3D-CMPs. As part of future work, we plan to explore support for variable process variations, and consider a service queue model that includes wait time of an application, for further improvements within our framework.

References

[1] V. B. Kleeberger, M. Barke, C. Werner, D. Schmitt-Landsiedel, and U. Schlichtmann, "A compact model for NBTI degradation and recovery under use-profile variations and its application to aging analysis of digital integrated circuits," *Microelectronics Rel.*, vol. 54, no. 6, pp. 1083–1089, 2014.

[2] D. Bergstrom, et al., "Intel's 45 nm CMOS technology," *Intel Technol. J.*, vol. 12, no. 2, pp. 131–144, Jun. 2008.

[3] E. Mintarno, et al., "Self-tuning for maximized lifetime energy-efficiency in the presence of circuit aging," *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.*, vol. 30, no. 5, pp. 760–773, May 2011.

[4] X. Huang, T. Yu, V. Sukharev and S. X.-D. Tan, "Physics-based electromigration assessment for power grid networks," in *Proc. Design Automation Conf.*, Jun. 2014, pp. 1–6.

[5] T. Okumura, F. Minami, K. Shimazaki, K. Kuwada, and M. Hashimoto, "Gate delay estimation in STA under dynamic power supply noise," *IEICE Trans. Fundamentals Electron. Commun. Comput. Sci.*, vol. 93, no. 12, pp. 2447–2455, 2010.

[6] N. H. Khan, S. M. Alam, and S. Hassoun, "System-level comparison of power delivery design for 2D and 3D ICs," in *Proc. IEEE Int. Conf. 3D Syst. Integr.*, 2009, pp. 1–7.

[7] W. Chan, A. B. Kahng, and S. Nath, "Methodology for electromigration signoff in the presence of adaptive voltage scaling," in *Proc. IEEE Int. Workshop Syst. Level Interconnect Prediction*, 2014, pp. 1–7.

[8] J. Lee and N. S. Kim, "Optimizing total power of many-core processors considering voltage scaling limit and process variations," in *Proc. ACM/IEEE Int. Symp. Low Power Electron. Des.*, 2009, pp. 201–206.

[9] S. Borkar, "Design perspectives on 22nm CMOS and beyond," in *Proc. Des. Automation Conf.*, 2009, pp. 93–94.

[10] J. Allred, S. Roy, and K. Chakraborty, "Designing for dark silicon: A methodological perspective on energy efficient systems," in *Proc. ACM/IEEE Int. Symp. Low Power Electron. Des.*, 2012, pp. 255–260.

[11] B. Raghunathan, Y. Turakhia, S. Garg, and D. Marculescu, "Cherry-picking: Exploiting process variations in dark-silicon homogeneous chip multi-processors," in *Proc. Conf. Des. Automation Test Europe*, 2013, pp. 39–44.

[12] S. S. Sapatnekar, "What happens when circuits grow old: Aging issues in CMOS design," in *Proc. Int. Symp. VLSI Technol. Syst. Appl.*, 2013, pp. 1–2.

[13] A. Tiwari and J. Torrellas, "Facelift: Hiding and slowing down aging in multicores," in *Proc. IEEE/ACM Int. Symp. Microarchitecture*, 2008, pp. 129–140.

[14] S. Feng, S. Gupta, A. Ansari, and S. Mahlke, "Maestro: Orchestrating lifetime reliability in chip multiprocessors," in *Proc. Int. Conf. High-Performance Embedded Archit. Compil.*, 2010, pp. 186–200.

[15] F. Paterna, A. Acquaviva, and L. Benini, "Aging-aware energy-efficient workload allocation for mobile multimedia platforms," *IEEE Trans. Parallel Distrib. Syst.*, vol. 24, no. 8, pp. 1489–1499, Aug. 2013.

[16] P. Mercati, A. Bartolini, F. Paterna, T. S. Rosing, and L. Benini, "Workload and user experience-aware dynamic reliability management in multicore processors," in *Proc. Design Autom. Conf.*, May 2013, Art. no. 2.

[17] J. B. Velamala, K. Sutaria, H. Shimizu, H. Awano, T. Sato, and Y. Cao, "Statistical aging under dynamic voltage scaling: A logarithmic model approach," in *Proc. IEEE Custom Integr. Circuits Conf.*, 2012, pp. 1–4.

[18] J. B. Velamala, K. Sutaria, T. Sato, and Y. Cao, "Physics matters: statistical aging prediction under trapping/detrapping," in *Proc. Design Automation Conf.*, Jun. 2012, pp. 139–144.

[19] J. B. Velamala, et al., "Compact modeling of statistical BTI under trapping/detrapping," *IEEE Trans. Electron Devices*, vol. 60, no. 11, pp. 3645–3654, Nov. 2013.

[20] V. Mishra and S. S. Sapatnekar, "The impact of electromigration in copper interconnects on power grid integrity," in *Proc. Des. Automation Conf.*, May 2013, Art. no. 88.

[21] S.-W. Chen, M.-H. Chang, W.-C. Hsieh, and W. Hwang, "Fully on-chip temperature, process, and voltage sensors," in *Proc. IEEE Int. Symp. Circuits Syst.*, 2010, pp. 897–900.
[22] P. Singh, E. Karl, D. Sylvester, and D. Blaauw, "Dynamic nbti management using a 45 nm multi-degradation sensor," *IEEE Trans. Circuits Syst. I: Regular Papers*, vol. 58, no. 9, pp. 2026–2037, Sep. 2011.
[23] A. Sassone, M. Petricca, M. Poncino, and E. Macii, "A fully standard-cell delay measurement circuit for timing variability detection," in *Proc. Int. Workshop Power Timing Model. Opt. Simul.*, 2013, pp. 239–242.
[24] S. Dighe, et al., "Within-die variation-aware dynamic-voltage-frequency-scaling with optimal core allocation and thread hopping for the 80-core teraflops processor," *IEEE J. Solid-State Circuits*, vol. 46, no. 1, pp. 184–193, Jan. 2011.
[25] S. R. Sarangi, B. Greskamp, R. Teodorescu, J. Nakano, A. Tiwari, and J. Torrellas, "VARIUS: A model of process variation and resulting timing errors for microarchitects," *IEEE Trans. Semiconductor Manufacturing*, vol. 21, no. 1, pp. 3–13, Jan. 2008.
[26] M. Arjomand and H. Sarbazi-Azad, "Voltage-frequency planning for thermal-aware, low-power design of regular 3-D NoCs," in *Proc. Int. Conf. VLSI Des.*, 2010, pp. 57–62.
[27] W. Jang, D. Ding, and D. Z. Pan, "A voltage-frequency island aware energy optimization framework for networks-on-chip," in *Proc. Int. Conf. Comput.-Aided Des.*, 2008, pp. 264–269.
[28] lp_solve 5.5.2.0. [Online]. Available: http://lpsolve.sourceforge.net/5.5/
[29] 3D-ICE open-source tool, [Online]. Available: http://esl.epfl.ch/3d-ice.html
[30] N. H. Khan, S. M. Alam, and S. Hassoun, "Power delivery design for 3-D ICs using different through-silicon via (TSV) technologies," *IEEE Trans. Very Large Scale Integr. Syst.*, vol. 19, no. 4, pp. 647–658, Apr. 2011.
[31] S. C. Woo, M. Ohara, E. Torrie, J. P. Singh, and A. Gupta, "The SPLASH-2 programs: Characterization and methodological considerations," *ACM SIGARCH Comput. Archit. News*, vol. 23, no. 2, pp. 24–36, 1995.
[32] C. Bienia, S. Kumar, J. P. Singh, and K. Li, "The PARSEC benchmark suite: Characterization and architectural implications," in *Proc. Int. Conf. Parallel Archit. Compilation Techn.*, 2008, pp. 72–81.
[33] T. E. Carlson, W. Heirman, and L. Eeckhout, "Sniper: Exploring the level of abstraction for scalable and accurate parallel multi-core simulation," in *Proc. Int. Conf. High Performance Comput. Netw. Storage Anal.*, 2011, Art. no. 52.
[34] S. Li, J. H. Ahn, R. D. Strong, J. B. Brockman, D. M. Tullsen, and N. P. Jouppi, "McPAT: An integrated power, area, and timing modeling framework for multicore and manycore architectures," in *Proc. IEEE/ACM Int. Symp. Microarchitecture*, 2009, pp. 479–480.
[35] K. Bhardwaj, K. Chakraborty, and S. Roy, "Towards graceful aging degradation in NoCs through an adaptive routing algorithm," in *Proc. ACM/EDAC/IEEE Des. Automation Conf.*, Jun. 2012, pp. 382–391.
[36] S. Pasricha and Y. Zou, "A low overhead fault tolerant routing scheme for 3D Networks-on-Chip," in *Proc. Int. Symp. Quality Electron. Des.*, 2011, pp. 1–8.
[37] M. Fattah, M. Daneshtalab, P. Liljeberg, and J. Plosila, "Smart hill climbing for agile dynamic mapping in many-core systems," in *Proc. Des. Autom. Conf.*, 2013, Art. no. 39.
[38] M. H. Haghbayan, A. Miele, A. M. Rahmani, P. Liljeberg, and H. Tenhunen, "A lifetime-aware runtime mapping approach for many-core systems in the dark silicon era," in *Proc. Des. Autom. Test Europe Conf. Exhib.*, 2016, pp. 854–857.
[39] H. Lin and Q. Xu, "Characterizing the lifetime reliability of many-core processors with core-level redundancy," in *Proc. Int. Conf. Comput.-Aided Des.*, 2010, pp. 680–685.
[40] D. Gnad, M. Shafique, F. Kriebel, S. Rehman, D. Sun, and J. Henkel, "Hayat: Harnessing dark silicon and variability for aging deceleration and balancing," in *Proc. IEEE Des. Autom. Conf.*, 2015, pp. 1–6.
[41] A. K. Singh, M. Shafique, A. Kumar, and J. Henkel, "Analysis and mapping for thermal and energy efficiency of 3-D video processing on 3-D multicore processors," *IEEE Trans. Very Large Scale Integr. Syst.*, vol. 24, no. 8, pp. 2745–2758, Aug. 2016.
[42] S. Rehman, F. Kriebel, D. Sun, M. Shafique, and J. Henkel, "dTune: Leveraging reliable code generation for adaptive dependability tuning under process variation and aging-induced effects," in *Proc. ACM Des. Autom. Conf.*, 2014, pp. 1–6.
[43] A. Das, A. Kumar, and B. Veeravalli, "Reliability and energy-aware mapping and scheduling of multimedia applications on multiprocessor systems," *IEEE Trans. Parallel Distrib. Syst.*, vol. 27, no. 3, pp. 869–884, Mar. 2016.
[44] A. Das, A. K. Singh, and A. Kumar, "Execution trace–driven energy-reliability optimization for multimedia MPSoCs," *ACM Trans. Reconfigurable Technol. Syst.*, vol. 8, no. 3, 2015, Art. no. 18.
[45] S.S. Sahoo, A. Kumar, and B. Veeravalli, "Design and evaluation of reliability-oriented task re-mapping in MPSoCs using time-series analysis of intermittent faults," in *Proc. IEEE Des. Autom. Test Europe Conf. Exhib.*, 2016, pp. 798–803.
[46] P. Singh, E. Karl, and D. Blaauw, "Compact degradation sensors for monitoring NBTI and Oxide degradation," *IEEE Trans. Very Large Scale Integrations Syst.*, vol. 20, no. 9, pp. 1645–1655, Sep. 2012.
[47] N. Kapadia and S. Pasricha, "VARSHA: Variation and reliability-aware application scheduling with adaptive parallelism in the dark-silicon era," in *Proc. IEEE/ACM Des. Autom. Test Europe. Conf. Exhib.*, Mar. 2015, pp. 1060–1065.
[48] S. Pasricha and Y. Zou, "A low overhead fault tolerant routing scheme for 3D networks-on-chip," in *Proc. IEEE Int. Symp. Quality Electron. Des.*, 2011, pp. 1–8.
[49] S. Pasricha and Y. Zou, "NS-FTR: A fault tolerant routing scheme for networks on chip with permanent and runtime intermittent faults," in *Proc. IEEE/ACM Asia South Pacific Des. Autom. Conf.*, 2011, pp. 443–448.

Venkata Yaswanth Raparti (S'15) received the bachelor's degree in electrical and electronics engineering from the Birla Institute of Technology and Science - Pilani, India, in 2011. From 2011-2013, he worked as a software engineer with Samsung Research Institute in Bangalore, India. He currently is working toward the PhD degree in the Electrical and Computer Engineering Department, Colorado State University, Fort Collins. His research interests include reliability aware task scheduling and mapping methodologies for 2D and 3D CMPs in Dark Silicon Era and networks-on- chip design for GPGPUs.

Nishit Kapadia (S'12) received the MS degree in electrical and computer engineering from California State University, Northridge, California, in 2005. He received the PhD degree in electrical and computer engineering from Colorado State University, in May 2016. From 2006 to 2008, he worked as a VLSI faculty in the Department of Electronics, Sardar Vallabhbhai National Institute of Technology, Surat, India. He currently works as a senior R&D engineer at Synopsys Inc. in Hillsboro, Oregon. His research interests are CAD for multi-core 2D and 3D Systems, networks-on-chip, power-, thermal-, variation-, and reliability-aware design methodologies, and VLSI Automation.

Sudeep Pasricha (M'02,-SM'13) received the BE degree in electronics and communication engineering from the Delhi Institute of Technology, India, in 2000, and the MS and PhD degrees in computer science from the University of California, Irvine, in 2005 and 2008, respectively. He is currently a monfort associate professor of electrical and computer engineering with Colorado State University, Fort Collins. His research interests include energy efficiency and fault tolerant design for network and memory architectures in multicore computing. He is or has been an organizing committee member and/or technical program committee member of various IEEE/ACM conferences such as DAC, DATE, CODES+ISSS, NOCS, ISQED, VLSID, and GLSVLSI. He is the recipient of several awards for his research contributions, including the 2015 IEEE TCSC Award for Excellence for a mid-career researcher and the 2013 AFOSR Young Investigator Award, as well as Best Paper Awards at the ACM SLIP 2016, ACM GLSVLSI 2015, IEEE AICCSA 2011, IEEE ISQED 2010, and ACM/IEEE ASPDAC 2006 conferences.