

# An Ageing-Aware and Temperature Mapping Algorithm for Multilevel Cache Nodes

EMMANUEL OFORI-ATTAH<sup>1</sup> AND MICHAEL OPOKU AGYEMAN<sup>1</sup>, (Senior Member, IEEE)

Faculty of Science and Technology, University of Northampton, Northampton NN1 5PH, U.K.

Corresponding author: Emmanuel Ofori-Attah (emmanuel.ofori-attah@northampton.ac.uk).

This work was supported by the Faculty of Arts, Science and Technology, University of Northampton.

**ABSTRACT** Increase in chip inactivity in the future threatens the performance of many-core systems and therefore, efficient techniques are required for continuous scaling of transistors. As of a result of this challenge, future proposed many-core system designs must consider the possibility of a 50% functioning chip per time as well maintaining performance. Fortunately, this 50% inactivity can be increased by managing the temperature of active nodes and the placement of the dark nodes to leverage a balance working chip whilst considering the lifetime of nodes. However, allocating dark nodes inefficiently can increase the temperature of the chip and increase the waiting time of applications. Consequently, due to stochastic application characteristics, a dynamic rescheduling technique is more desirable compared to fixed design mapping. In this paper, we propose an Ageing Before Temperature Electromigration-Aware, Negative Bias Temperature Instability (NBTI) & Time-dependent Dielectric Breakdown (TDDB) Neighbour Allocation (ABENA 2.0), a dynamic rescheduling management system which considers the ageing and temperature before mapping applications. ABENA also considers the location of active and dark nodes and migrate task based on the characteristics of the nodes. Our proposed algorithm employ Dynamic Voltage Frequency Scaling (DVFS) to reduce the Voltage and Frequency (VF) of the nodes. Results show that, our proposed methods improve the ageing of nodes compared to a conventional round-robin management system by 10% in temperature, and 10% ageing.

**INDEX TERMS** Dark-silicon, power management, run-time mapping, many-core systems, dynamic voltage frequency scaling (DVFS), task migration, power states, dark core.

## I. INTRODUCTION

Excessive thermal rise in many-core systems caused by uncontrolled leakage power as transistor size decreases leads to the Dark-Silicon challenge, where a section of many-core chips have to powered-off or under-clocked to function within a power budget; whilst considering the temperature threshold as well as meeting application demands [1]–[4]. Thermal Design Power (TDP) and Thermal Safe Power (TSP) [5] are examples of power budgeting techniques used to prevent chip damage with the latter offering more power for resource activity. Consequently, during run-time, DVFS, Power Gating (PG), Task Migration (TM) and Near Threshold Computing (NTC) are used to control the power.

However, inefficient exploitation of Dark-Silicon can lead to a poverty-stricken performance system. This is because, different resource activation patterns result in different thermal profiles and thus, presents an opportunity

The associate editor coordinating the review of this manuscript and approving it for publication was Ahmed A. Zaki Diab<sup>1</sup>.

for performance optimization. An optimized activation of resources can determine the maximum allowable power budget.

Contiguous resource activation offers better application performance however, the heat produced from active working components leads to thermal hotspots and limits the number of activated resources. This leads to more power being consumed. Additionally, this can also threaten the reliability and lifetime of the chip.

On the other hand, non-contiguous resource activation can degrade the performance of the system without respect for applications which require inter communication. As a matter of fact, the communication overhead caused by the distance between nodes which require inter communication would require processors running faster which eventually leads to more power being consumed [6]. Therefore, to leverage the disparity between performance and temperature, the non-active nodes, which we refer to as “dark” or “dumso” interchangeably are placed in-between active nodes. Over the years, various techniques have been proposed which exploit

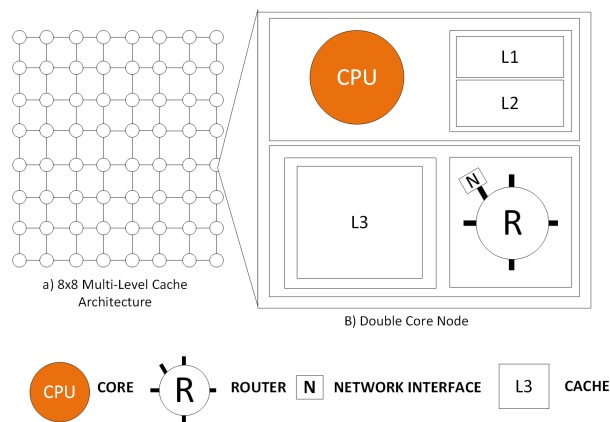


FIGURE 1. Multi-level-cache architecture.

the use of dark nodes to improve the performance of active cores. An active node can function at a higher frequency when dark nodes are placed near it. However, over use of dark nodes can lead to the starvation of newly arrive applications as a result of insufficient nodes in a large management systems [7].

Moreover, this not only causes a performance challenge, but it also introduces expensive packaging and cooling costs coupled with reliability and ageing of resources with the latter being vastly ignored by many proposed techniques. As a matter fact, ITRS predict that as technology scales further down, future chips will age faster [8].

To overcome these challenges, several design aspects have been proposed in literature [9]–[13], [13], [14]. One design aspect which can be exploited in different ways is architectural heterogeneity: whether this being the implementation of different sizes of resources as depicted in [15] or incorporating CPU/GPU [16], [17]. Yang *et al.* [15] proposed the Quad-core cluster. The Quad-core cluster consists of four different types of cores (High Performance, General Purpose, Power Saving, Low Energy). Based on the application demands, the appropriate core is selected for computation. Additionally, DVFS and TM are employed to scale down VF and migrate tasks to cores at run-time [18]–[20].

Another technique widely used is the efficient use of management systems [21], [22]. Others on the other hand, reduce power consumption in Network-on-Chip (NoC) Components [23]–[27] for an optimized performance. Fig. 1 depicts a Multi-Level Cache Node comprised of core, router and a Multi-Level Cache Architecture (MCA). Although MCA is mentioned in this work, we only address the power consumption of the nodes. We do not extensively address power related issues as this will be done in future work. In this work, we seek to improve the temperature and lifetime of nodes by using the lifetime of nodes as the main factor to allocate applications and tasks. The highlights of the proposed scheme are as follows:

This paper addresses the challenges by proposing a Dark-Silicon patterning approach that handles repeated

workloads where the temperature and lifetime of nodes are considered. The applications and tasks are assumed as fixed. The main contributions of this paper can be summarised as:

- We present two Dark-Silicon algorithms to improve the lifetime reliability of many-core systems.
- The first algorithm improves the lifetime of nodes by migrating tasks from hot nodes to dark nodes at run-time.
- The second algorithm monitors and improve the lifetime of nodes by migrate tasks after an epoch. Additionally, the algorithm monitors the Lifetime of two neighbouring nodes and then selects the node with the highest MTTF to improve the Lifetime.
- Additionally, the proposed method reduces temperature by ensuring an active node is surrounded by dark nodes to allow it to function at higher frequency.
- The proposed approach is compared with the conventional round-robin against the following parameters: Temperature, Hotspot, Lifetime and Utilization.
- Results from the experiment conducted shows how optimised the proposed method is.

The paper is organized as follows: Section II briefly discusses related work about heterogeneous nodes, dynamic application mapping which considers computation and communication intensive application. Section III presents an observation into the round-robin mapping algorithm and in Section IV, both proposed approaches are presented. Section V presents the experimental results. Finally, Section VI concludes the paper and discusses future work.

## II. RELATED WORK

The allocation of resources in many-core systems have been the focus of technology since the emergence of the Dark-Silicon phenomenon. However, [16] states that, previous task-resource allocation only considers mapping applications contiguously without regard for generated heat amongst resources. Active components generate heat and over time spread among neighbouring resources causing thermal hotspot. Consequently, over time, this affects the lifetime reliability of the resources. This causes accelerated EM, NBTI & TDDDB and leads to function slower than the other [28].

Work that considers heat dissipation does not emphasise the lifetime reliability of systems and count node utilization as a crucial metric when optimising performance.

Xiaohang Wang *et al.* [7] proposed a virtual mapping algorithm to estimate the number of dark cores required for an application to efficiently execute. The aim of the virtual mapping algorithm is to prevent inefficient use of dark cores to ensure there are enough free cores for incoming applications. This algorithm considers communication and computation applications. However, node utilization is not considered as a metric when choosing the first node.

Kanduri *et al.* [29] presented adBoost, a thermal aware performance boosting system which boost the performance of active cores by the efficient use of dark cores. The algorithm

employed in this system maps applications spatially to avoid hot spot. First Node selection is selected based on a finding a node that is far away using MapPro from an active node with sufficient nodes around it for mapping of application task. Unfortunately, node utilization is not used as a property when selecting the first node.

The following work, however, addresses this challenge by considering computation and communication demands when activating resources.

Reza *et al.* [16] proposed a resource management system for heterogeneous NoCs which examines the performance requirement of an application (Communication or Computation Intensive) before distributing its tasks amongst the appropriate resources. In this management system, the chip is partitioned into clusters with a set of CPUs and GPUs. The architecture present employs a CPU from each cluster and assigns them as Cluster managers to monitor and configure the resources within that cluster. This information is then feed-backed to a global manager which uses the BalancedMap mapping application algorithm to activate the feasible resources based on the application. The BalancedMap algorithm maps communication intensive applications to CPUs and computation intensive application to GPUs.

The global manager is assigned by calculating the CPU with the shortest distance from all the assigned Cluster Managers. The BalancedMap Mapping Algorithm works by classifying the application into two groups (Communication and Communication Intensive). During CPU and GPU node selection, the node with the lowest peak energy is selected. If more than one node appears to meet the demand, the node with the lowest utilization is selected to improve the thermal hotspot. Power consumption is minimized by configuring the link, node voltage and the shutting down of idle routers when links are not carrying any communication traffic. TM is employed when a node being utilized exceeds it warning threshold.

Rahmani *et al.* [30] considers reliability in Dark-Silicon systems by employing reliability-aware power allocator to monitor the ageing of nodes and to reduce the amount of workload on stressed out regions that are experiencing fast ageing. Consequently, the voltage/Frequency of the nodes are scaled down to prevent the nodes from functioning at full throttle.

Rathore *et al.* [31] on the other hand, proposed a HiMap, a hierarchical mapping approach which reduces the lifetime of nodes by mapping application to healthier nodes whilst also placing dark nodes amongst the selected region. This mapping approach considers Process Variation when assessing the ageing and reliability of the node. The assessment is done by checking nodes for process variation, temperature, and ageing. Weaker nodes are then used as dark nodes.

Mohammed *et al.* [32] proposed a technique which uses both TS and DVFS for an optimise system performance. TM is used to move cores from active to dark cores. In this architecture, the dark and active cores are all running

concurrently and therefore makes it easier for task to be swap amongst them. Unfortunately, this increases the temperature and thus, DVFS is applied.

Rathore *et al.* [33] proposed a technique to optimise the lifetime reliability for many-core systems. Similar to our proposed technique, they employ clusters. Although their proposed system maps tasks to the node, hotspot can still occur because of mapping tasks to high MTF nodes in the same area. Alternatively, Ansari *et al.* [34] proposed a Thermal-Aware Standby-Sparing Technique (TASS) in a heterogeneous multi-core system. The multi-core system consists of two pairs of cores: High Performance Core (HP) and Low Performance Core (LP). The main tasks of applications are scheduled on the HP whilst the backup tasks are scheduled on the back up cores. Unfortunately, there is no mention of high temperature between active components which are next to each other. This can still hotspot. In our purposed method, we compare between neighbours to reduce hotspot.

In our previous work [35], we presented architectural saving techniques to improve the power efficiency for many-core systems. Based on the study conducted, it can be deduced the fraction of powered-off nodes in Dark-Silicon chips can be improved by power budgeting, architectural heterogeneity, NoC interconnect, Cache Memory, and run-time management. Particularly, resource allocation using application mapping through run-time management methods.

Similar to the above-mentioned techniques, we approach the Dark-Silicon challenge by permitting only 50% of the nodes available to function at full through per time. However, unlike previous proposed approaches that statically and randomly generate dark and active nodes, our proposed algorithm gracefully selects dark and active nodes based on their current lifetime and temperature.

### III. PROBLEM STATEMENT

At any given time,  $A$  number of applications dynamically arrive in the system. Each application spawns several  $T$  threads.  $\{A_0, A_1, A_2, A_3, \dots, A_{A-1}\}$ . The objective is to allocate each application to several  $N$  nodes  $\{N_0, N_1, N_2, N_3, \dots, N_{N-1}\}$  to meet the application deadline without aggravating the temperature of the chip. Fortunately, the overall temperature depends on several factors: The number of actively working nodes, the position and the voltage frequency of nodes also have an impact. Therefore, the heat generated can be model as:

$$H = \sum \mathcal{P}^i \times \tau^{amb} \times \tau^i$$

where  $\mathcal{P}^i$  is the power dissipated by the node,  $\tau^{amb}$  is the temperature of neighbouring nodes and  $\tau^i$  is the ambient temperature.

Contiguous mapping techniques spawn application threads next to each other without considering the temperature and heat generated by neighbouring nodes. The heat dissipated by active nodes haphazardly affects neighbouring nodes and pushes nodes to reach their critical temperatures even when

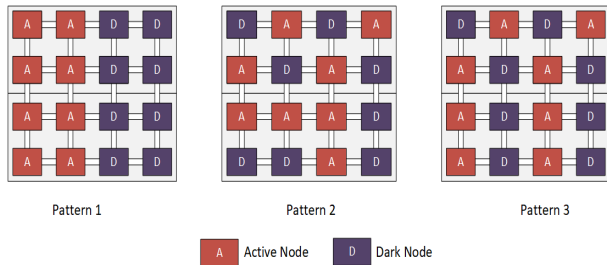


FIGURE 2. Dark-silicon patterns.

they are operating under low V/F. Pattern 1 in Fig.2 is an example of such an approach. Non-contiguous mapping avoids this challenge by dispersing application and threads across the chip. Unfortunately, dispersing tasks across the chip increases the communication latency between threads which require inter communication and thus, increases the power consumption and temperature.

A naive mapping of two applications can cause the temperature of the chip to be aggravated. Fig. 2 depicts three different patterns. Pattern 1 is likely to result in a high temperature chip because all the nodes that have been selected for mapping are in the same region. Additionally, incoming applications will be continuously mapped to the same active nodes. Pattern 2 and 3 on the other hand have dispersed applications across the chip. This offers low temperature; however, consideration of node performance must be considered when making such a decision to prevent bottlenecks. This can either be, the temperature of the node, the lifetime of the node, the location of the node or the effect that a mapping of a thread to a node can have on an incoming application.

Therefore, different mapping approaches results in various thermal profile of the chip. An optimal approach can allow up to 10 applications to be mapped without aggravating the temperature of the many-core system. Hence, our proposed approach.

#### IV. PROPOSED APPROACH

This section presents the proposed technique. The proposed technique aims to optimise the performance of the many-core systems under several thermal constraints. Due to modern processors supporting several low-power states termed C-states (C0, C1, C2, ...) [36], we employ DVFS to manage the frequencies and to shut down parts of the nodes based on their thermal profile.

Fig. 3 shows the overall architecture of the considered system. The platform contains a set of homogeneous processing tiles organised in a 2-D mesh-based topology and connected by an NoC.

##### A. PROPOSED FRAMEWORK

The proposed many-core platform is partitioned into several clusters of equal sizes in an  $N_X \times N_X$  grid of MCA tiles.

The clusters are formed by specifying the number of nodes in the system, the type of mesh required, the application with

the highest number of tasks and the number of tasks. 3 shows a 64-node system with  $4 \times 4$  sized clusters. To determine the number of nodes per cluster in a 64-node system, the application with the highest number of tasks is considered. If an application has 8 tasks, then four clusters will be created with 16 nodes (8 active and 8 dark). This is because ABENA employs the 50% dumso rule dark-silicon.

Each cluster has been allocated a module to monitor the temperature and lifetime of nodes and to utilise DTM techniques to the nodes accordingly. Each module then returns the average temperature and lifetime of its cluster to a centralised resource manager. Based on the information collected, applications are assigned to the right cluster. Additionally each node has per-core DVFS capability [37].

##### 1) PROPOSED ALGORITHM

Majority of dynamic mapping applications do not consider ageing even though some nodes age faster than others. To improve performance, temperature and power budget are usually used as measurements or parameters for enabling and disabling nodes. Over time, permanent faults caused by Electromigration (EM), Negative Dependent Temperature Instability (NBTI), and Time Dependent Dielectric Breakdown (TDDB) reduces the lifetime and increases the systems Mean Time To Failure (MTTF).

Therefore, efficient techniques are required to improve the lifetime of nodes as well maintaining the performance. Hence why we proposed a new method which uses the MTTF as main option for choosing dumso and active nodes. The proposed approach is not limited to just one phenomenon, it can be applied to other ageing mechanisms. However, for the purpose of this study, we consider TDDB. The MTTF of a node or system is an immediate value generated under existing circumstances. This can change depending on the circumstance, parameters, and application characteristics. In contrast, Aging is accumulated over a period of time.

The proposed algorithm (ABENA) is detailed in Algorithm 1. ABENA Monitors the following parameters at run-time: the temperature of all the clusters  $T_c$ , the average MTTF of all clusters  $MTTF_c$ , the temperature of all active nodes  $A = \{a_0, \dots, a_0\}$ , the temperature of all dark nodes  $D = \{d_0, \dots, d_0\}$ , the MTTF of all nodes  $MTTF_n$ , and the frequencies of all nodes  $F$ . Given  $n \times n$  number of nodes, ABENA divides the many-core system into several  $C_{SET}$ . We employ this Dark-Silicon cluster mapping approach because it allows flexibility between restricting application threads to a specific area whilst also allowing the tasks to be dispersed in that region. For each  $C$ , we calculate and compare the average  $MTTF$  by monitoring all  $MTTF$  of the  $N$  in each cluster. ABENA selects the first  $C$  for mapping by calculating the average  $MTTF$  of nodes from each  $C$ . The highest  $MTTF$  cluster is selected as the prime candidate for mapping. Additionally, if two clusters return the same  $MTTF$ , the  $C$  with the lowest  $T$  is chosen (line 2-5). In line 5 - 6, active nodes from each region are chosen based on their  $T$ . Step 4 in Fig. 3 depicts this approach.



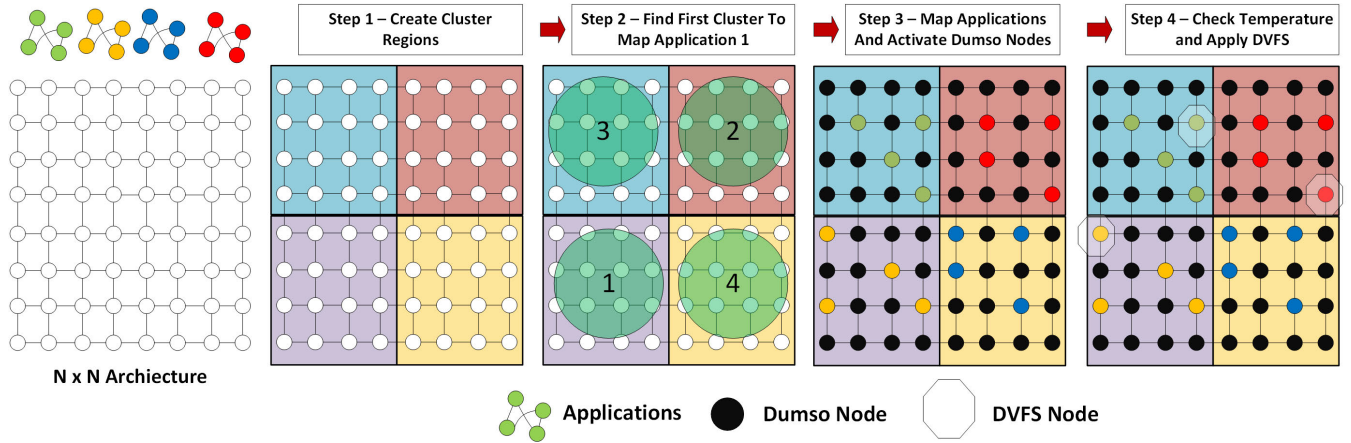


FIGURE 3. Proposed architecture.

**Algorithm 1** ABENA Algorithm

```

Data:  $C_1 \dots C_N$ 
Result: Low to High Clusters
1  $N$  : Nodes  $C$ :Clusters ;
2  $C_{SET}$  : ClusterSet ;
3 ReturnAverage : ReturnAverageoftheCluster ;
4  $MTTF^o$  : MeanTimeToFailure ;
5  $T_{TRS}$  : TemperatureThreshold ;
6 Temp : Temperature ;
7  $N_{Frequency}$  : Temperature ;
8 for  $C \in C_{SET}$  do
9 | ReturnAverage  $\leftarrow N_K \div (\sum N_1^{MTTF} \times N_K^{MTTF})$ ;
10 end
11 FirstCluster  $\leftarrow C$ withtheHighestMTTFAverage
12 * Initial Mapping * Select 50% of the nodes with the
lowest temperature nodes from each cluster
13 if  $N_{Temp^o} > T_{TRS}$  then
14 |  $N_k \leftarrow ApplyDVFS N_{Frequency} - 10\%$ 
15 else
16 end
    
```

Unfortunately, this mapping approach is unpredictable. In a large many-core system, hotspot can easily occur when four neighbouring nodes with the lowest  $T$  is selected. To reduce this ABENA implements a Neighbour-Temperature-Aware Node-Allocation (NANA) (Algorithm 2) which forms clusters between two neighbouring nodes.

This ensures that,  $A$  nodes are always surrounded by the  $D$  to hotspot. The active node between the two neighbouring nodes is selected based on the node with the Highest MTTF. (Line 1). If the active node's temperature is more than the  $T_{TRS}$ , the  $F$  is decreased. The proposed design is implemented with different voltage levels. In the experiment conducted, the frequency is reduced by 10%. The 10% is not a special value that guarantees the temperature violations. From a design perspective, reducing the frequency by 1% will not have any significant effect and thus we use 10%. Additionally, if the temperature of the nodes does not change, the task is migrated

**Algorithm 2** NANA Algorithm

```

Data:  $NC_1 \dots C_N$  Node Cluster
1 foreach  $NC_1 \in C_N$  do
2 | if  $N_1 > N_2$  then
3 | |  $N_k \leftarrow AssignThread$  if  $N_{1Temp^o} > T_{TRS}$  : then
4 | | |  $N_k \leftarrow ApplyDVFS N_{Frequency} - 10\%$ 
5 | | else
6 | | else
7 end
    
```

TABLE 1. The used symbols description.

Symbol	Description
$N$	Nodes
$A_{SET}$	Set of all active nodes
$A$	Active node
$D_{SET}$	Set of all dark nodes
$D$	Dark node
$C_{SET}$	Set of all Clusters
$C$	Cluster
$F$	Frequency
$T$	Temperature
$T_c$	Set of transient temperature of all clusters
$T_n$	Set of transient temperature of all nodes
$T_{TRS}$	Temperature Threshold
$MTTF_c$	Average MTTF of a cluster
$MTTF_n$	Average MTTF of nodes

unto the neighbouring dumso node. Although, ABENA aims to keep 32 nodes active, if all applications have been mapped, the idle active nodes are shut down to allow the working nodes to function at a higher frequency. ABENA assumes that a source node goes to C1 State at the initial phase of migration maintaining the contents of its L1/L2 cache.

**V. PERIODIC WORKLOADS**

Algorithm 3 depicts the algorithm ABENA-P, the periodic version of ABENA for mapping applications. ABENA-P checks for the MTTF of neighbouring nodes and assign tasks to nodes with the highest MTTF (Line 3). After assigning the tasks, ABENA-P then checks the temperature of the node and reduce the frequency of the node until it meets the tem-

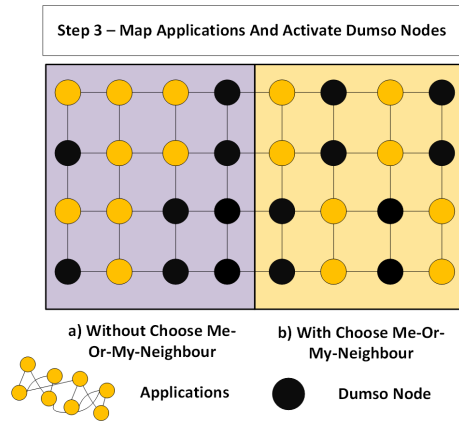


FIGURE 4. 4 × 4 architecture with the proposed ageing.

perature threshold. When ABENA-P is applied, intervention happens once every epoch, which is in months or hours resulting in minimal or no performance penalty.

**Algorithm 3** ABENA Periodic

```

Data: Epoch1 . . . . . EpochN : Epoch
Data: N : Nodes
Data: NC1 . . . NC1 : Neighbouring Clusters
Data: D: Dark node
1 foreach Epoch1 ∈ EpochN do
2   if N1,MTTF < N2,MTTF then
3     N1 ← D if N1,Tempo > TTRS : then
4       Nk ← ApplyDVFS NFrequency - 10%
5     else
6   else
7     N2 ← D if N1,Tempo > TTRS : then
8       Nk ← ApplyDVFS NFrequency - 10%
9     else
10    end
11 end
    
```

Consequently, to accommodate more applications, ABENA utilises the Flow Node Algorithm. This enables an application which is not heavily dependent on inter-communication to be mapped to nodes in different clusters. Similar to [31], at the end of every each epoch, applications are mapped to healthy nodes.

**VI. ENVIRONMENTAL EVALUATION**

This section presents the experimental setup, results, and discussions of the proposed work.

**A. EXPERIMENTAL SETUP**

Experiments and evaluations are conducted using an extended version of Sniper [38] called LifeSim [39], which integrates McPat [40] for power simulation, HotSpot [41] to model the thermal profile and generate thermal values. Additionally, LifeSim also integrates RAMP [42] to model the lifetime of the nodes. There are other power models such as Wattch [43]




FIGURE 5. Normalised comparison of the average peak temperatures of the chip by running the two algorithms on a 4.0 GHz frequency.




FIGURE 6. Temperature of ABENA.




FIGURE 7. Temperature of RR.

and PowerTrain [44] which can also be used to estimate the cores’ power profiles. However, to the best of our knowledge, McPat and HotSpot are the most widely tools used in literature to measure and assess temperature and power.

RAMP [42] is a dynamic reliability management tool which calculates the expected lifetime of cores based on the current temperature and its utilization. RAMP divides the processor into a few structures – ALUs, FPUs, register files, branch predictor, caches, load-store queue, instruction window – and applies the analytic models to each structure as an aggregate. Our proposed architecture uses the lifetime of nodes generated by RAMP as the parameter for assigning applications.

LifeSim simulator consists of a set of applications which spawn several threads which run simultaneously and can

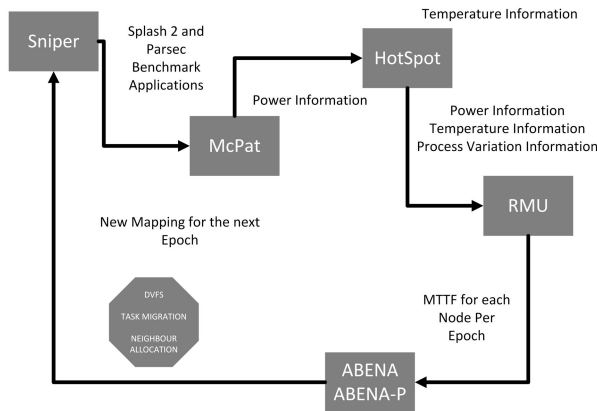


FIGURE 8. The environmental setup of the proposed method.

be picked from any of the available benchmarks. When the simulation starts, application threads are spawned on nodes specified by a mapping file. The experimental setup of the proposed work is shown in 8. During the simulation, LifeSim uses the previously mentioned integrated tools to generate the following values for analysis: McPAT to generate power values, HotSpot to generate thermal values using the MCPAT generated power values and RAMP to estimate the lifetime of the nodes.

LifeSim provides two modes of run-time. These are preemptive and run-time non-preemptive scheduling. For this this particularly case study, we use non-preemptive scheduling. Non-preemptive scheduling decisions are made ahead of time and therefore does not delay mapping. Non-preemptive also employs epoch schedule. Unlike run-preemptive scheduling where intervention occurs every specified second, non-preemptive scheduling occurs after every epoch which is usually in months and hours and thus time overhead and performance is very minimal. Similar to [45] we could not measure the switching overhead. Sniper-sim does not simulate the operating system and the context switch managed by it but only the user-space. In theory, context switching usually takes 5-7n but in this case study, intervention occurs after every epoch [33]. The profiling results (power, temperature and MTTF) of are logged in parallel to the simulation, each sampling interval. We invoke our proposed method during this process to compare and analyse the results being generated to assign applications, reduce the frequency of nodes, and migrate tasks among nodes. We compared the proposed method (ABENA-P) with a Dark-Silicon round-robin mapping algorithm with realistic computation applications (radix, oceans, fit and barnes from Splash-2 benchmark [46]) to verify and validate ABENA. By default, ABENA uses the default length epoch (1 Month).

Table 2 depicts the parameters for the simulated configuration. In our previous paper [47], we used a 45nm processor to evaluate ABENA in a 4 × 4 architecture. To prove that it can be applied to large scale architectures, we evaluated ABENA on many-core system that consists of 64 nodes, where 32 are active and 32 are dumso. The nodes are connected using an

TABLE 2. Summary of simulation parameters.

Parameter	Value
Nodes	64
Core Frequency	2.0 - 4.0 GHZ
Core Size	5.71 x 5.71
L1 Data Cache	32KB, 4-way set-associative, 64B Block Size
L1 Instruction Cache	32KB, 4-way set-associative, 64B Block Size
L2 Cache	512KB, 4-way set-associative, 64B Block Size
L3 Cache	8129, 23 Block Size
Applications	Radix, oceans, fft, barnes
Technology	22 nm
Network size	8 X 8 NoC
Temperature threshold	80 °
Failure Mechanism	TDDB
Algorithm °	
ABENA-P	

TABLE 3. Workload characteristics.

Parameter	Value
Temperature threshold	80°
Parameters	Temperature, Average MTTF, Utilisation

8 × 8 mesh NoC with 22nm processor. Each node consists of a private 32 KB L1 data cache, 32 KB L1 instruction and a 512KB L2 Cache.

To get accurate readings, we modified sniper to support Dark-Silicon and conducted the experiment 10 times. We compared our proposed method with the conventional Round-Robin (RR). This comparison was conducted into two stages. At the first stage, we conducted a study on both algorithms with 4.0 GHZ frequency to examine the worst-case scenarios; Primarily, this was done to test if the activation of nodes in a specific area affects the overall temperature. The threshold temperature was set to 80°C. The second stage experiment was conducted at a frequency of 2.0 GHZ. In this phase of the experiment, the DVFS in ABENA was turned-off. The parameters that the results are conducted on are temperature, MTTF, Hotspot and Utilization. In this research, the failure mechanism that is used is TDDB. Wu *et al.* concluded in [48] that temperature and voltage affect the lifetime of ultra-thin gate oxide due to TDDB. As a result of this, the MTTF due to TDDB,  $MTTF_{TDDB}$ , at a temperature,  $T$ , and a voltage,  $V$ , can be model as:

$$MTTF_{EM} \propto \left(\frac{1}{V}\right)^a (e^{-bT}) e^{\frac{(X + Y/T + ZT)}{kT}}$$

where the values  $a = 78, b = -0.081, X = 0.759ev, Y = -66.8evK, and Z = -8.081 - 4ev/K$  are fitting parameters [48].

### B. COMPARATIVE RESULTS AND ANALYSIS

Table 3 shows the workload characteristics. For a wide range of results, we compared both approaches with two frequencies (2.0 & 4.0 GHZ)

#### 1) EVALUATION OF TEMPERATURE BETWEEN ABENA AND RR USING 4.0 GHZ FREQUENCY

Fig. 5 shows a normalised temperature comparison between ABENA and RR. The RR architecture achieves a 3% advantage over ABENA because 50% of its nodes are not utilized.




FIGURE 9. Normalised comparison of the average MTF between ABENA and RR using Frequency 4.0 GHz.




FIGURE 10. Normalised comparison of the average utilization of the chip by running the two algorithms.

In contrast, the highest node from both algorithms shows that RR has the highest temperature caused by having active nodes functioning next to each other. The active nodes in ABENA achieves a 30% advantage of RR. In addition, ABENA reduces hotspot in its architectures. Fig. 7 is an image of the average temperature of ABENA and RR. By placing dark nodes next to each active node and using DVFS, we minimise hotspot and keep the temperature of the nodes below the temperature cap. This is reflected in the highest temperature of ABENA being 79°C compared to 82°C with RR. ABENA achieves a better temperature performance for all the active nodes.

2) EVALUATION OF MTF AND UTILISATION BETWEEN ABENA AND RR USING A 4.0 GHZ FREQUENCY

Fig. 9 shows a normalised comparison of the MTF from simulated results between ABENA and RR. ABENA outperforms RR by 20% when comparing the average of active nodes. As previously mentioned, due to half of nodes in RR being inactive, the average of RR is higher than ABENA.

However, ABENA outperforms RR when comparing the node with the lowest and highest active MTF. By using the lifetime as a parameter, the lifetime of a node is evaluated before a task is assigned to it.

Fig. 10 shows the utilization comparison between both approaches. It can be concluded from Fig. 10 that the gap between the highest node and lowest from ABENA is smaller compared to RR. This proves how architectures




FIGURE 11. Normalised comparison of the average peak temperatures of the chip by running the two algorithms.




FIGURE 12. Temperature of ABENA.




FIGURE 13. Temperature of RR.

which implement RR suffers perform it is implemented. There is 30% reduction in this gap when ABENA is applied.

Moreover, based on the results presented, it is obvious that ABENA outperforms the conventional round-robin. ABENA ensures that the nodes are functioning below the temperature cap to prevent thermal issues whilst improving the lifetime of nodes. However, we also notice that the introduction of DVFS in ABENA provides an advantage for ABENA. Therefore, we conducted another set of experiments where the frequency was reduced to 2.0 without the use of DVFS.

3) EVALUATION OF TEMPERATURE BETWEEN ABENA AND RR USING 2.0 GHZ FREQUENCY

Fig. 11 shows the average temperature of the two methods when they are evaluated with a frequency of 2.0. One can see that the average temperature of all the algorithms is below the temperature threshold but once again RR achieves the






FIGURE 14. Normalised comparison of the average MTFF temperatures between ABENA and RR using Frequency 2.0.

TABLE 4. Comparison between ABENA and Round-Robin.

Parameter	Round-Robin	ABENA
Hotspot	High	Low
Lifetime	Medium	Low
Temperature	Medium	Low

lowest average temperature. For the rest of the results, it can be concluded that ABENA outperforms RR in every area. The highest node and lowest node from the study revealed that ABENA minimises the temperature by 10% and 15% respectively. Consequently, ABENA outperforms RR by 20% when comparing the average of active nodes. Additionally, the highest node and lowest node from ABENA also outperforms RR by 10% respectively. This proves that even without the DVFS, ABENA achieves a better temperature performance because of the pattern at which active and dark nodes are selected. Fig. 11 shows the temperature of ABENA. Similar to Fig. 5, the temperature is evenly spread across the chip. There is evenly spread of temperature across the chip.

4) EVALUATION OF MTFF AND UTILISATION BETWEEN ABENA AND RR USING A 2.0 GHZ FREQUENCY

Fig. 14 compares the MTFF of the two methods. Similar to the results from the frequency 4.0 GHZ, the highest node from both algorithms is RR. However, ABENA outperforms RR when it comes to the highest and lowest MTFF node by 10%. This further highlights the disparity between the highest node and the lowest. Additionally, ABENA outperforms RR by 10% in average of active nodes, 15% and 10% in highest and lowest active node respectively.

Fig. 15 shows the idle percentage of nodes used in both algorithms. From the image, it is obvious that the idle percentage of nodes in the round-robin algorithm is high and presents a major issue for the lifetime of nodes. As a matter of the fact, there is 40% gap between the highest and lowest node from RR. In contrast, the highest node and lowest node from ABENA is a 20% gap. Additionally, ABENA outperforms RR in every area when comparing active node utilization.

VII. SUMMARY OF CONTRIBUTION

Moreover, it can be concluded that ABENA outperforms RR in hotspot minimisation, temperature, MTFF and utilization.




FIGURE 15. Normalised comparison of the average utilization of the chip by running the two algorithms.

Fig. 4 summaries the results. By activating a dark node next to dark node, ABENA improves the temperature performance of many-core systems by 10%; Primarily, the active nodes. In addition, ABENA also improves the lifetime of nodes by more 5% by efficiently utilization nodes. Throughout the study, ABENA bridged the disparity between the lowest and highest utilized. ABENA achieved this by using the life-time of nodes as an additional and main parameter before assigning tasks.

VIII. CONCLUSION AND FUTURE WORK

In this paper, we presented ABENA, a run-time lifetime algorithm to optimise the lifetime of nodes whilst improving the temperature and reducing hotspot. The temperature, MTFF and utilization under two different frequencies were evaluated and compared to the traditional round-robin algorithm. The results obtained show that, ABENA outperforms the conventional algorithm with or without DVFS. Although results show that the average temperature of both approaches for MTFF are similar, the average of active nodes shows that ABENA improves the average lifetime of all active nodes by 20%. This is evident by comparing the highest MTFF Node of both approaches. With more stress being placed on more nodes, the lifetime of the chip will diminish faster. ABENA delivers a balanced chip. Furthermore, this effectively improves the hotspot performance of the chip. This is because, by ensuring that there is always at least one dark node next to an active node, the temperature reduces. In contrast, the conventional round-robin algorithm creates hotspot which exceeds the temperature cap even though half the nodes of the chip are dark. Additionally, this shows that with a larger scale network and an increased frequency, the round-robin algorithm will cause thermal issues. In our future work, we will focus on task migration in every epoch to efficiently migrate tasks among active and dark nodes in a single workload. We will also evaluate the impact of consistent task migration.

REFERENCES

[1] D.-C. Juan, S. Garg, J. Park, and D. Marculescu, "Learning the optimal operating point for many-core systems with extended range voltage/frequency scaling," in *Proc. Int. Conf. Hardw./Softw. Codesign Syst. Synth. (CODES+ISSS)*, Sep. 2013, pp. 1–10.

- [2] A. Kanduri, M.-H. Haghbayan, A.-M. Rahmani, P. Liljeberg, A. Jantsch, and H. Tenhunen, "Dark silicon aware runtime mapping for many-core systems: A patterning approach," in *Proc. 33rd IEEE Int. Conf. Comput. Design (ICCD)*, Oct. 2015, pp. 573–580.
- [3] T. S. Muthukaruppan, M. Pricopi, V. Venkataramani, T. Mitra, and S. Vishin, "Hierarchical power management for asymmetric multi-core in dark silicon era," in *Proc. 50th Annu. Design Autom. Conf. (DAC)*, May 2013, pp. 1–9.
- [4] M. Shafique, S. Garg, T. Mitra, S. Parameswaran, and J. Henkel, "Dark silicon as a challenge for hardware/software co-design: Invited special session paper," in *Proc. Int. Conf. Hardw./Softw. Codesign Syst. Synth.*, Oct. 2014, pp. 1–10.
- [5] S. Pagani, H. Khdr, J.-J. Chen, M. Shafique, M. Li, and J. Henkel, "Thermal safe power (TSP): Efficient power budgeting for heterogeneous manycore systems in dark silicon," *IEEE Trans. Comput.*, vol. 66, no. 1, pp. 147–162, Jan. 2017.
- [6] W. Liu, L. Yang, W. Jiang, L. Feng, N. Guan, W. Zhang, and N. Dutt, "Thermal-aware task mapping on dynamically reconfigurable network-on-chip based multiprocessor system-on-chip," *IEEE Trans. Comput.*, vol. 67, no. 12, pp. 1818–1834, Dec. 2018.
- [7] X. Wang, A. K. Singh, B. Li, Y. Yang, H. Li, and T. Mak, "Bubble budgeting: Throughput optimization for dynamic workloads by exploiting dark cores in many core systems," *IEEE Trans. Comput.*, vol. 67, no. 2, pp. 178–192, Feb. 2018.
- [8] (2014). *International Technology Roadmap for Semiconductors (ITRS)*. [Online]. Available: <http://public.itrs.net>
- [9] E. Ofori-Attah and M. O. Agyeman, "A survey of recent contributions on low power NoC architectures," in *Proc. Comput. Conf.*, Jul. 2017, pp. 1086–1090.
- [10] E. Ofori-Attah and M. O. Agyeman, *A Survey of Low Power NoC Design Techniques*. New York, NY, USA: Association for Computing Machinery, 2017.
- [11] E. Ofori-Attah, X. Wang, and M. O. Agyeman, "A survey of low power design techniques for last level caches," in *Proc. ARC*, 2018, pp. 217–228.
- [12] E. Ofori-Attah, W. Bhebhe, and M. O. Agyeman, "Architectural techniques for improving the power consumption of NoC-based CMPs: A case study of cache and network layer," *J. Low Power Electron. Appl.*, vol. 7, no. 2, p. 14, 2017.
- [13] I. Hoxha and M. O. Agyeman, "High performance heterogeneous multi-core architectures: A study," in *Proc. 3rd Int. Symp. Comput. Sci. Intell. Control*, Sep. 2019, pp. 1–5.
- [14] E. Ofori-Attah and M. O. Agyeman, "A survey of low power NoC design techniques," in *Proc. 2nd Int. Workshop Adv. Interconnect Solutions Technol. Emerg. Comput. Syst.*, Jan. 2017, pp. 22–27.
- [15] L. Yang, W. Liu, N. Guan, M. Li, P. Chen, and E. H. M. Sha, "Dark silicon-aware hardware-software collaborated design for heterogeneous many-core systems," in *Proc. 22nd Asia South Pacific Design Autom. Conf. (ASP-DAC)*, Jan. 2017, pp. 494–499.
- [16] M. F. Reza, D. Zhao, and M. Bayoumi, "Power-thermal aware balanced task-resource co-allocation in heterogeneous many CPU-GPU cores NoC in dark silicon era," in *Proc. 31st IEEE Int. Syst. Chip Conf. (SOCC)*, Sep. 2018, pp. 260–265.
- [17] Y. Zhang, L. Peng, X. Fu, and Y. Hu, "Lighting the dark silicon by exploiting heterogeneity on future processors," in *Proc. 50th Annu. Design Autom. Conf. (DAC)*, 2013, pp. 1–7.
- [18] I. Anagnostopoulos, V. Tsoutsouras, A. Bartzas, and D. Soudris, "Distributed run-time resource management for malleable applications on many-core platforms," in *Proc. 50th Annu. Design Autom. Conf. (DAC)*, 2013, pp. 1–6.
- [19] E. L. D. S. Carvalho, N. L. Calazans, and F. G. Moraes, "Dynamic task mapping for MPSoCs," *IEEE Design Test Comput.*, vol. 27, no. 5, pp. 26–35, Sep. 2010.
- [20] J. Castrillon, A. Tretter, R. Leupers, and G. Ascheid, "Communication-aware mapping of KPN applications onto heterogeneous MPSoCs," in *Proc. 49th Annu. Design Autom. Conf. (DAC)*, Jun. 2012, pp. 1262–1267.
- [21] T. S. Muthukaruppan, M. Pricopi, V. Venkataramani, T. Mitra, and S. Vishin, "Hierarchical power management for asymmetric multi-core in dark silicon era," in *Proc. 50th Annu. Design Autom. Conf. (DAC)*, 2013, pp. 1–9.
- [22] 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. Design, Autom. Test Eur. Conf. Exhib. (DATE)*, Mar. 2016, pp. 854–857.
- [23] C. Feng, Z. Liao, Z. Lu, A. Jantsch, and Z. Zhao, "Performance analysis of on-chip bufferless router with multi-ejection ports," in *Proc. IEEE 11th Int. Conf. ASIC (ASICON)*, Nov. 2015, pp. 1–4.
- [24] C. Fallin, C. Craik, and O. Mutlu, "CHIPPER: A low-complexity bufferless deflection router," in *Proc. IEEE 17th Int. Symp. High Perform. Comput. Archit.*, Feb. 2011, pp. 144–155.
- [25] B. K. Daya, L. S. Peh, and A. P. Chandrakasan, "Towards high-performance bufferless NoCs with sceptor," *IEEE Comput. Archit. Lett.*, vol. 15, no. 1, pp. 62–65, Jan./Jun. 2016.
- [26] C. Li and P. Ampadu, "A compact low-power eDRAM-based NoC buffer," in *Proc. IEEE/ACM Int. Symp. Low Power Electron. Design (ISLPED)*, Jul. 2015, pp. 116–121.
- [27] N. Nasirian and M. Bayoumi, "Low-latency power-efficient adaptive router design for network-on-chip," in *Proc. 28th IEEE Int. Syst. Chip Conf. (SOCC)*, Sep. 2015, pp. 287–291.
- [28] K. A. Bowman, S. G. Duvall, and J. D. Meindl, "Impact of die-to-die and within-die parameter fluctuations on the maximum clock frequency distribution for gigascale integration," *IEEE J. Solid-State Circuits*, vol. 37, no. 2, pp. 183–190, Aug. 2002.
- [29] A. Kanduri, M.-H. Haghbayan, A. M. Rahmani, M. Shafique, A. Jantsch, and P. Liljeberg, "AdBoost: Thermal aware performance boosting through dark silicon patterning," *IEEE Trans. Comput.*, vol. 67, no. 8, pp. 1062–1077, Aug. 2018.
- [30] A. M. Rahmani, M.-H. Haghbayan, A. Miele, P. Liljeberg, A. Jantsch, and H. Tenhunen, "Reliability-aware runtime power management for many-core systems in the dark silicon era," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 25, no. 2, pp. 427–440, Feb. 2017.
- [31] V. Rathore, V. Chaturvedi, A. K. Singh, T. Srikanthan, R. Rohith, S.-K. Lam, and M. Shafique, "HiMap: A hierarchical mapping approach for enhancing lifetime reliability of dark silicon manycore systems," in *Proc. Design, Autom. Test Eur. Conf. Exhib. (DATE)*, Mar. 2018, pp. 991–996.
- [32] M. S. Mohammed, A. K. Al-Dhamari, A. A.-H. A. Rahman, N. Paraman, A. A. M. Al-Kubati, and M. N. Marsono, "Temperature-aware task scheduling for dark silicon many-core system-on-chip," in *Proc. 8th Int. Conf. Modeling Simulation Appl. Optim. (ICMSAO)*, Apr. 2019, pp. 1–5.
- [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 Perform. Comput., Netw., Storage Anal. (SC)*, 2011, pp. 1–12.
- [34] M. Ansari, S. Safari, S. Yari-Karin, P. Gohari-Nazari, H. Khdr, M. Shafique, J. Henkel, and A. Ejlali, "Thermal-aware standby-sparing technique on heterogeneous real-time embedded systems," *IEEE Trans. Emerg. Topics Comput.*, early access, Oct. 20, 2021, doi: [10.1109/TETC.2021.3120084](https://doi.org/10.1109/TETC.2021.3120084).
- [35] E. Ofori-Attah, X. Wang, and M. Agyeman, "A survey of system level power management schemes in the dark-silicon era for many-core architectures," *EAI Endorsed Trans. Ind. Netw. Intell. Syst.*, vol. 5, no. 15, Sep. 2018, Art. no. 155569.
- [36] *Advanced Configuration and Power Interface (ACPI) Specification*, Forum U.E.F.I., Beaverton, OR, USA, 2020.
- [37] W. Kim, M. S. Gupta, G.-Y. Wei, and D. Brooks, "System level analysis of fast, per-core DVFS using on-chip switching regulators," in *Proc. IEEE 14th Int. Symp. High Perform. Comput. Archit.*, Feb. 2008, pp. 123–134.
- [38] T. E. Carlson, W. Heirman, S. Eyermer, I. Hur, and L. Eeckhout, "An evaluation of high-level mechanistic core models," *ACM Trans. Archit. Code Optim.*, vol. 11, no. 3, pp. 1–25, Oct. 2014.
- [39] R. Rohith, V. Rathore, V. Chaturvedi, A. K. Singh, S. Thambipillai, and S.-K. Lam, "LifeSim: A lifetime reliability simulator for manycore systems," in *Proc. IEEE 8th Annu. Comput. Commun. Workshop Conf. (CCWC)*, Jan. 2018, pp. 375–381.
- [40] 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. 42nd Annu. IEEE/ACM Int. Symp. Microarchitecture (Micro)*, 2009, pp. 469–480.
- [41] R. Zhang, M. R. Stan, and K. Skadron, "Hotspot 6.0: Validation, acceleration and extension," Univ. Virginia, Charlottesville, VA, USA, Tech. Rep., 2015.
- [42] J. Srinivasan, S. V. Adve, P. Bose, and J. A. Rivers, "The case for lifetime reliability-aware microprocessors," in *Proc. 31st Annu. Int. Symp. Comput. Archit.*, 2004, p. 276.
- [43] D. Brooks, V. Tiwari, and M. Martonosi, "Watch: A framework for architectural-level power analysis and optimizations," *ACM SIGARCH Comput. Archit. News*, vol. 28, no. 2, pp. 83–94, 2000.

- [44] W. Lee, Y. Kim, J. H. Ryoo, D. Sunwoo, A. Gerstlauer, and L. K. John, "PowerTrain: A learning-based calibration of McPAT power models," in *Proc. IEEE/ACM Int. Symp. Low Power Electron. Design (ISLPED)*, Jul. 2015, pp. 189–194.
- [45] V. Rathore, V. Chaturvedi, A. K. Singh, T. Srikanthan, and M. Shafique, "Longevity framework: Leveraging online integrated aging-aware hierarchical mapping and VF-selection for lifetime reliability optimization in manycore processors," *IEEE Trans. Comput.*, vol. 70, no. 7, pp. 1106–1119, Jul. 2021.
- [46] 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.*, 1995, pp. 24–36.
- [47] E. Ofori-Attah and M. O. Agyeman, *ABENA: An Ageing before Temperature Electromigration-Aware Neighbour Allocation for Many-Core Architectures*. New York, NY, USA: Association for Computing Machinery, 2021, pp. 146–152.
- [48] J. Srinivasan, S. V. Adve, P. Bose, and J. A. Rivers, "The case for lifetime reliability-aware microprocessors," in *Proc. 31st Annu. Int. Symp. Comput. Archit.*, 2004, pp. 276–287.



**EMMANUEL OFORI-ATTAH** received the B.Eng. degree in computer systems engineering from the University of Northampton, where he is currently pursuing the Ph.D. degree. He has extensive knowledge in programming and has industrial experience as a Full Stack Developer. Previously, he worked as a Graduate Teaching Assistant with the Faculty of Science and Technology, from 2013 to 2018. He is a Lecturer with the University of Northampton. He is currently the author of three journals with over six conference publications since his inception to field. His research interests include high performance computing and embedded systems [VLSI SoC design, computer architecture, re-configurable computing, embedded systems, and the Internet-of-Things (IoT)]. He has won several awards. These include the University of Northampton's 3 Minute Thesis (3MT) Competition (People's Award and Judges Awards) and an award from the Worshipful Company of Information Technologists (WCIT).



**MICHAEL OPOKU AGYEMAN** (Senior Member, IEEE) received the Ph.D. degree from Glasgow Caledonian University, U.K. He is currently a Professor and the Program Leader of computer systems engineering at the University of Northampton (UoN), U.K. He represents the Research Community of UoN at University Senate. He is the Postgraduate (PGR) Lead for the Faculty of Arts Science and Technology and co-chairs the University's PGR Supervisory Forum. He has over ten years experience in embedded systems engineering. Previously, he was with Intel Embedded System Research Group, The Chinese University of Hong Kong (CUHK), as a Research Fellow. He is the author of five books, two book chapters, and over 70 publications in major journals and conference proceedings. His main research interests include across three main strands and disciplines such as embedded systems high performance computing [VLSI SoC design, computer architecture, reconfigurable computing, wired and wireless NoCs, smart rehabilitation solutions, embedded systems, and the Internet-of-Things (IoT)], business administration (neuromarketing, advertising, and market research), and pedagogy. He is a Chartered Engineer (C.Eng.) of the IET, the Chartered Manager (C.Mgr.) of CMI, and a fellow of the Higher Education Academy, U.K. He is a Technical Committee Member of several conferences such as IEEE ICCSN, IEEE ICBD, and IEEE ICCT, and serves as a reviewer of several conferences and journals including IEEE Access. His work on wireless NoC has attracted two best paper awards in IEEE/IFIP EUC 2016 and Euromicro DSD 2016. He was a recipient of the 2018 International Changemaker of the year award in the first U.K. Ashoka U Changemaker Campus. He has been a Guest Editor of the *EAI Endorsed Transactions on Industrial Networks and Intelligent Systems*.

• • •