Reliability-Aware Runtime Power Management for Many-Core Systems in the Dark Silicon Era

Amir M. Rahmani, *Member, IEEE,* M. Hashem Haghbayan, *Student Member, IEEE,*
Antonio Miele, *Member, IEEE,* Pasi Liljeberg, *Member, IEEE,* Axel Jantsch, *Member, IEEE,*
Hannu Tenhunen, *Member, IEEE*

Abstract—Power management of networked many-core systems with runtime application mapping becomes more challenging in the dark silicon era. It necessitates to consider network characteristics at runtime to achieve better performance while honoring the peak power upper bound. On the other hand, power management has a direct effect on chip temperature which is the main driver of the aging effects. Therefore, alongside performance fulfillment, the controlling mechanism must also consider the current cores' reliability in its actuator manipulation to enhance the overall system lifetime in the long term. In this paper, we propose a multi-objective dynamic power management technique that uses current power consumption and other network characteristics including the reliability of the cores as the feedback while utilizing fine-grained voltage and frequency scaling and per-core power gating as the actuators. In addition, disturbance rejecter and reliability balancer are designed to help the controller to better smooth power consumption in the short-term and reliability in the long-term, respectively. Simulations of dynamic workloads and mixed criticality application profiles show that our method not only is effective in honoring the power budget while considerably boosting the system throughput, but also increases the overall system lifetime by minimizing aging effects by means of power consumption balancing.

Index Terms—Dark Silicon; Power Management; Feedback Controller; Networks-on-Chip; Runtime Mapping; Lifetime Reliability

I. INTRODUCTION

THE number of transistors per chip still steadily increases by about $2.0\times$ for each technology node generation. However, due to increased power densities and consequent thermal issues, power budget represents a constraint in the usability of such computational resources: in fact, the number of transistors that can be used at the same time on a chip increases with a lower trend (only $1.4\times$ each generation [1]). Consequently, only a subset of the available processing cores can be used at full throttle at any of time, while the rest of the resources have to be left inactive, thus representing Dark Silicon [2], [3]. Thermal issues do not represent a problem only in the short term by idling resources, but also in the long term. In fact, as discussed in the ITRS reports in 2011 [4]

Amir M. Rahmani, M. Hashem Haghbayan, Pasi Liljeberg, and Hannu Tenhunen are with Department of Information Technology, University of Turku, Turku, Finland, 20520. Antonio Miele is with Dipartimento di Elettronica, Politecnico di Milano, Milano, Italy, 20133. Axel Jantsch is with TU Wien, Austria.
E-mail: amirah@utu.fi, mohhag@utu.fi, pasi.liljeberg@utu.fi, axel.jantsch@tuwien.ac.at, hannu@kth.se
Hannu Tenhunen is also with the Department of Industrial and Medical Electronics, Royal Institute of Technology (KTH), Kista, Sweden, 16440.

due to transistor shrinking, the higher temperatures cause the devices to be more susceptible to aging and wear-out phenomena (such as time dependent dielectric breakdown, thermal cycling, or electromigration), thus decreasing the system lifetime. As stated in [5], many failure mechanisms are exponentially dependent on temperature, and a 10-$15°C$ difference in operating temperature may result in a $2\times$ difference in the overall lifespan of a device.

An investigated solution to mitigate dark silicon phenomenon is the Near-Threshold Computing (NTC), also dubbed as Dim Silicon [6], and more in general acting on the voltage/frequency (VF) levels of the processing cores. The basic idea is to increase the number of concurrently active cores by considerably lowering their operating VF level. However, to implement an efficient NTC-based approach, an intelligent and stable power management mechanism is required. Such a mechanism becomes more challenging, considering that current and future many-core systems integrate various hundreds of networked resources and use to execute a highly variable and unpredictable workload. Therefore, to accurately handle an upper limit on power consumption (fixed Thermal Design Power, TDP or dynamic Thermal Safe Power, TSP [7]) and at the same time guarantee a certain level of Quality of Service in the running applications, a feedback control mechanism, monitoring several parameters of the system is mandatory. Finally, the quest of such a feedback control mechanism is even more motivated when we aim also at limiting the aging trend and to prolong the overall system lifetime.

The previous work on feedback-based dynamic power management for multi-/many-core systems can be classified into two main categories: i) techniques which use workload and network characteristics as feedback (e.g., queue utilization and injection rate), and then adjust voltage/frequency of processing elements, routers, or voltage/frequency islands (VFI) accordingly (e.g., [8], [9]), and ii) power budgeting (i.e., capping) techniques [10], [11] which utilize chip/per-core power measurement and per-core performance counters (i.e., core utilization) as feedback, and then apply Dynamic Voltage and Frequency Scaling (DVFS) or per-core power gating (PCPG) techniques to optimize the system performance within a fixed power cap (i.e., TDP). Even though all the techniques in these categories efficiently save and control the power consumption for their target platforms, they are not comprehensive and do not take aging effects into account. The techniques in the first category do not consider any safe upper bound on the total system power consumption (i.e., TDP) at

runtime, and therefore, they do not feed any power metric back to the management unit. The power capping techniques from the second category are mainly targeted for bus-based systems. Therefore, they are unable to address the power management issues in many-core systems which are typically based on a Network-on-Chip (NoC) and with multiple applications running simultaneously. Finally, very few approaches in the literature consider also aging issues in the dynamic mapping decisions, with [11] as a notable exception. In conclusion we contend that *dark silicon awareness necessitates an efficient multi-objective feedback-based control approach which considers many parameters such as workload characteristics, per-core power, aging and performance measurements, network-load, and total chip power consumption all together.*

We propose a novel comprehensive dark silicon-aware and reliability-aware power management framework for NoC-based many-core systems under limited power budget (both TDP and TSP) and running dynamic workloads, by supporting runtime mapping. This framework benefits from a multi-objective feedback-based controller which i) pursues performance and lifetime reliability optimization while fulfilling the power budget, ii) acts on PCPG and DVFS, and iii) considers a large set of monitoring parameters such as workload characteristics, network congestion, and power/performance/aging characteristics of processing cores. A preliminary version of the approach has been proposed in [12]. We here extend it to consider also lifetime reliability issues together with power and performance optimization, as follows:

- Adding the reliability analysis unit to calculate fine grained reliability based on the temperature feedback profile from the system.
- Developing novel decision policies targeted for two different operating modes: Over-boosting Mode, when the system is experiencing an intensive workload, and Reliability-aware Mode, when the non-intensive workload offers the controller the opportunity to prolong the system lifetime.
- Extending the metrics to performing VF scaling to consider reliability of the system.
- Adding an additional reliability balancing module running at coarse time intervals.
- Evaluating the efficiency of our approach to provide high performance while prolonging the system's lifetime and fulfilling the given power budget.

The rest of the paper is organized as follows. In Section II, related work is presented. An overall view of the NoC-based many-core system together with the envisioned runtime management framework is presented in Section III while subsequent Sections IV and V present the internals of the proposed power controller and companion monitoring infrastructure. A experimental evaluation of the proposed approach is provided in Section VI, and, finally, Section VII draws conclusions.

II. RELATED WORK

Various approaches for dark-silicon-aware power management have been proposed in the literature. In [10], a hierarchical power management controller for asymmetric multi-

cores is demonstrated on an ARM big.LITTLE platform. Ma *et al.* [11] propose a similar technique, called *PGCapping*, by exploiting power gating and DVFS for power capping in symmetric multi-core processors. These two approaches use a feedback loop to manage power consumption, and their main limitation is that they target a bus-based architecture. Therefore, they do not consider many relevant issues manifesting in NoC-based many-core systems. In [13] the authors propose a method to allocate shared energy-efficient accelerators among competing applications to minimize the overall power while fulfilling the applications' deadlines. In [14], a hierarchical technique is proposed to partition TDP among running applications. In [15], the authors propose a power management technique for many-core systems. In their platform applications use feedback to monitor other running applications' activity to self-manage their power consumption and degree of parallelism. However, none of these works consider any power feedback from the system at runtime and instead use a pre-defined dedicated power budget.

Chen *et al.* [16] present a power allocation technique for many-core system performance improvement under power constraints. They formulate a performance optimization problem and apply an optimal power allocation method using on-line distributed reinforcement learning. This contribution does not consider on-chip communication. However, it should be noted that this work is orthogonal to our multi-objective control management and can complement and enhance our method by integrating the concept of on-line learning towards more efficient global power budget reallocation. In [8], a control based approach is proposed to minimize dynamic power in MPSoC made of multiple, voltage frequency islands (VFIs). Their goal is to determine optimal operating frequencies for both PEs and routers. This work is not dark silicon aware either, as they do not utilize feedback from power sensors to avoid violating the TDP/TSP.

Haghbayan *et al.* [17] present a power management technique for many-core systems using power feedback from the system to meet the TDP bound. This technique is also categorized into the class of single objective control approaches as it lacks feedbacks from workload characteristics and per-core performance measurements from the system during DVFS. In addition, this technique is designed for fixed TDP and does not benefit from a dedicated disturbance rejector to handle sudden overshoots when new applications commence execution.

A more advanced approach for power management of many-core systems in the dark silicon era has been proposed in [12]. The controller implements a feedback control loop collecting a large set of information from the running workload and the underlying architecture (such as network congestion and performance/power consumption status of the various cores), and acting on both PCPG and DVFS to optimize system performance while fulfilling the available power budget.

With the exception of the bus-based power capping technique presented in [11], none of the above approaches take into account the effects on the thermal stress with its impact on aging. AS a matter of fact, reliability management cannot be performed by only considering TDP or TSP; in fact, TDP/TSP can only guarantee that excessive temperature peaks are not
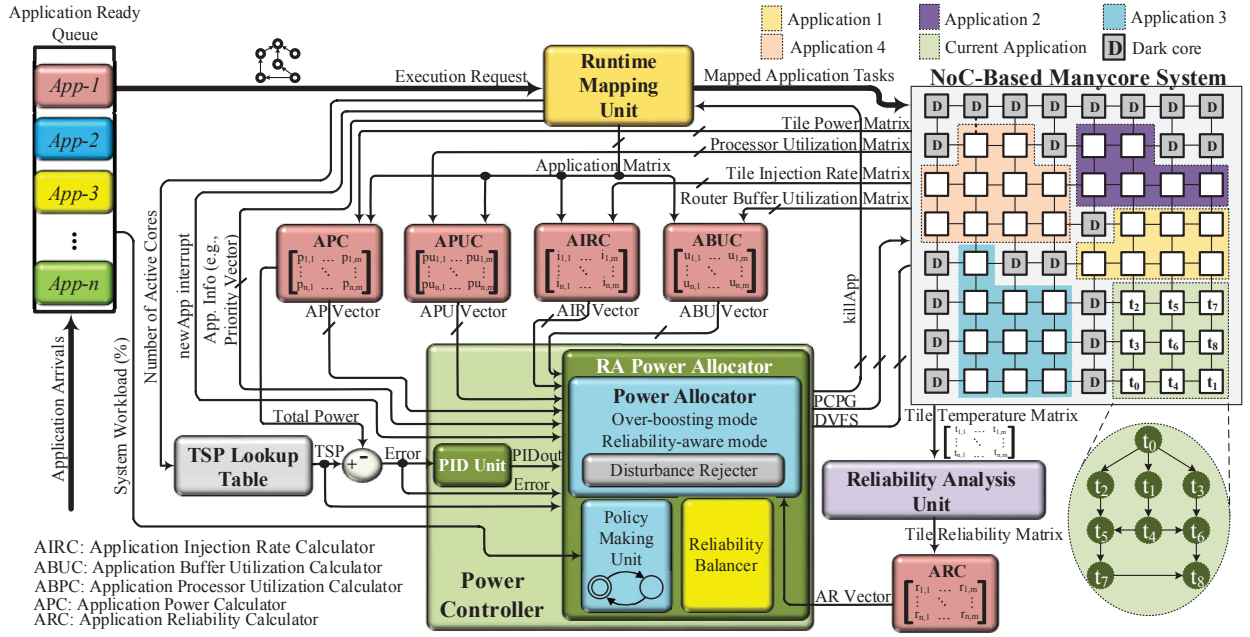
Figure 1: Overview of our multi-objective dark silicon aware power management system

reached but does not mitigate the effects of high temperatures in the long term. For this reason some approaches [18], [19] act on DVFS to balance the trade-off among performance, power consumption and reliability. However, they are mainly targeted for bus-based systems.

III. SYSTEM OVERVIEW

Figure 1 shows the overall architecture of the considered system. It consists of a classical hardware many-core platform and the above software stack, composed of the newly-proposed runtime resource management layer, and the running applications' workload. The hardware platform we consider captures the main characteristics of modern many-core systems, such as the Intel Single-chip Cloud Computer (SCC, [20]) or the Kalray MPPA [21]. This platform contains a set of homogeneous processing tiles organized in a 2D mesh-based topology and connected by a Network-on-Chip. Moreover, the platform is connected to a host machine controlling all the activities.

Each tile contains a processing core provided with a local private instruction/data memory and a communication controller connecting to a corresponding router in the NoC; communications among tiles are performed by means of a message-passing protocol over the NoC. Finally, we assume each tile to be provided with HW sensors for temperature and power consumption measurements (as in [11], [16], [20]), performance counters measuring the number of executed instructions (as in [20]) and actuators for PCPG and DVFS (as in [11], [16]). It is worth noting that the proposed approach can be easily applied to variants of the considered architecture having different network topologies or actuation knobs (e.g., coarser-grained VF at cluster level).

The many-core system generally executes data-intensive and highly-parallel applications, organized in a group of inter-dependent processing tasks. Applications may also present priority requirements (e.g., soft realtime, non-realtime) correlating to the level of expected Quality of Service (QoS). Such requirements are here expressed in terms of a priority index that can be used to rank applications while the QoS in terms of a minimum throughput (or a deadline on the overall execution time) to be guaranteed. To be executed, each application needs to be mapped on the architecture, by reserving a group of tiles on which the application's tasks are dispatched.

The system supports the concurrent execution of a highly dynamic workload, characterizing the nowadays on-demand computing scenario. Multiple applications dynamically arrive with an unknown trend and are characterized by different structures (in terms of number of tasks and their organization), latencies and QoS requirements. Therefore, the architecture loads a software layer for runtime resource management performing two main tasks: runtime mapping and dynamic power management. More precisely, similar to Intel SCC or Kalray MPPA, the main part of this software layer, implementing all the decision policies for mapping and power management, is executed on the host machine. Then, each tile loads a minimal scheduler running the received tasks and managing transmissions. Furthermore, the tile scheduler executes also a set of routines performing monitoring activities by means of the available sensors.

The Runtime Mapping Unit (RMU) is in charge of performing the dispatching of the workload. Entering applications are initially stored in a ready queue, and the RMU performs the mapping of a ready application according to the availability of idle resources (i.e., the tiles set as dark since not allocated for the execution of any other applications) and power, assigned by the power controller. The RMU also categorizes information of the applications running on the system (defined as running applications' information, *RAI*) such as the idle/active tile matrix, the application mapping

matrix, and the applications' priority list.

The second component within the runtime resource management layer is the Power Controller. This module is in charge of accurately handling the available power budget and distributing it among the various groups of tiles executing the running applications. The available power budget is used as a target value for Power Controller Unit as shown in Figure 1 and can be TDP or TSP (only TSP is mentioned in the figure). The activity of the controller is supported by means of a set of auxiliary monitoring modules pre-processing *RAI* gathered from the RMU and raw measures from the architectural sensors, and providing a set of high level information and statistics driving the decisions on DVFS and PCPG actuation. It is worth noting that the power controller does not scale the voltage and frequency (VF) of on-chip interconnection network components (e.g., NoC routers and links), to ensure that there are no unnecessary delays in the transmission thus causing additional static power consumption of waiting tile. Finally, the controller provides also commands to the RMU to coordinate mapping decisions with the power management.

The novel contribution of this paper is the Power Controller, together with the companion utility modules, while, a state-of-the-art RMU [22] will be considered. The monitoring and power management policies are presented in the following.

IV. System Monitoring

As depicted in Figure 1, a set of auxiliary modules perform monitoring activities on the system to collect various kinds of information. Each module is implemented by means of two parts: 1) a software routines periodically running on each tile that collect data from HW sensors, and send them to the host machine, and 2) the counterpart function integrated in the software layer on the host machine that will group all values of the same type in a matrix-based data structure.

Many aspects in the current system status may be relevant to take the most efficient choices on DVFS and PCPG actuation in order to get high performance for each running application, and, at the same time, optimize the power consumption of the system and slowing down the aging trend of the processing units. As an example, it is intuitive that the change of the current VF level of a set of cores allocated to the execution of an application will have a direct effect on both the performance and the power consumption. Nevertheless, it is also fundamental to monitor the traffic on the NoC as congestion in the network may represent a bottleneck that would make the increase of VF on the tiles to have no gain on the performance, while adding an additional contribution on power consumption. Such a scenario becomes even more complex when considering that the proposed approach aims at considering also, as a novelty, system's lifetime reliability as the third objective. As a conclusion, the auxiliary modules collect detailed information on the status of the processing units, i.e. the tiles, and of the communication infrastructure at router level, and aggregate these data with the granularity of each single application. The auxiliary modules integrated in the proposed framework are discussed in this section.

Application-level Power Calculator (APC). This module periodically samples the tile's power consumption by means

of the available HW sensors and sends it to the host machine. We assume this sensed value to contain power contributions of all components within the tile, i.e. the processing core, the memory, the communication controller and the connected router. Then, the host machine groups all these values in a Tile Power Matrix. According to the Application Matrix provided by the RMU, the APC module calculates the current power consumption of each application by aggregating the values of all tiles involved in its execution creating the Application Power Vector (*APV*), and transmits it to the Power Controller. As depicted in Figure 1, the input matrices (as the ones of the subsequent modules), which store the collected tile-level measures, have the same shape and dimension of the architecture grid. Then, the output is a vector with a length equal to the number of running applications.

Application-level Processor Utilization Calculator (APUC). As in [11], the APC exploits the performance counters integrated within each tile to collect the Processor Utilization Matrix, which measures the utilization of each processing unit (in percentages) during the previous timing window. Then, the APUC module aggregates the utilization values of the various tiles allocated for the same application by means of a sum. Such aggregated metrics, showing the performance demand of the running applications, form the Application Processor Utilization Vector (*APUV*), passed to the Power Controller.

Application-level Packets Injection Rate Calculator (AIRC). Another useful information to decide how to tune the VF level of the various tiles is related to the amount of traffic generated by each single application. As noted in [23], a communication-intensive application may not benefit from an increase of the VF level in terms of higher performance. Indeed, the network would represent a performance bottleneck and an increase in the VF level would only worsen such the overall network congestion and consume more power without any benefit. For this reason, we consider applications' network intensity in order to classify them into intensive and non-intensive categories in the power management process, and we use application packets injection rate as a metric that closely correlates to network intensity.

The packets injection rate of each task is measured at network interface of the allocated tile based on its recent history. The AIRC receives such measures in the form of a Tile Injection Rate Matrix, and it accordingly calculates the average injection rate for each application by selecting the related values according to the Application Matrix. The obtained Application-level Packets Injection Rate Vector (i.e., $AIRV$) is transmitted to the Power Controller.

Application-level Buffer Utilization Calculator (ABUC). Monitoring the utilization level of the routers involved in the execution of a single application (i.e., the routers connected to the allocated tiles) is also fundamental to detect possible traffic congestions that would affect the application performance. Indeed, the network capacity does not grow proportionally to accommodate traffic generated with increasing number of tiles [24]. As a consequence, the network traffic gets more easily congested proportionally with the number of tiles, and the overall throughput per-tile decreases. Therefore, the

total network performance gives a diminishing return due to increased communication distance. Figure 2 shows that this issue leads to a network performance gap where every tile is not able to send or receive packets in every cycle. Thus, the Power Controller needs to monitor and detect such network performance gaps to make efficient decisions on power management actuation.

In the considered platform, each router is equipped with a buffer utilization meter (as in [8]) monitoring router congestion levels in its recent history. More precisely, it dynamically measures the traffic in terms of average congestion level C_t at any given cycle t, by calculating the moving average of packet flow in every link of a router. Then, similarly to the previous modules, the ABUC elaborates the Router Buffer Utilization Matrix and the Application Matrix to compute the Application Buffer Utilization Vector (*ABUV*) where buffer utilizations are averaged per application. As a final note, AIRC and ABUC are complementary in the analysis of the network status. In fact, the former monitors the amount of traffic generated by each single application, that depends on the characteristics of the application and the VF configuration of the processing cores, while the latter represents the actual congestion of the network due to the packets flow among the various routers, that depends on the intrinsic characteristics of the network in terms of speed and routing scheme that would cause the actual accumulation of packets in the buffers.

TSP Lookup Table. The last input for the Power Controller is the available power budget. TDP has been classically used to define such a power budget at design time in order to avoid excessive temperatures. However as shown in recent works [7], TDP leads to a suboptimal usage of the device and results in large performance loss. For this reason, in [7] a dynamic power budgeting approach, called Thermal Safe Power (TSP), is defined as a function of the number of active (i.e., non-dark) tiles in a system. In this work, we use TSP_{worst} strategy which defines the safe power budget by considering the worst-case mapping scenario (i.e., assuming all the active tiles are physically packed and influencing the temperature of their adjacent tiles). From a practical point of view, our approach is based on the pre-computation of the TSP_{worst} values for different number of active tiles in the system to build a TSP lookup table. Thus, at runtime, the *number of active tiles* can be used as an input of this lookup table to get the corresponding safe power budget P_{TSP}^{worst}. The TSP lookup table is queried whenever an application enters or leaves the

Figure 2: The share of network bandwidth per-core diminishes with increasing number of cores as shown here for uniform traffic.

system at runtime since these two events may cause a change in the number of active tiles. It should be noted that if using a fixed TDP value is desired, the lookup table can be simply replaced with the fixed value.

Reliability Monitor and Application-level Reliability Calculator (ARC). The framework is provided also with a monitor devoted to the observation of the aging status of the various tiles, The most precise way to monitor the aging of each component is by using hardware wear-out sensors. However, even though there is a large effort on the study and the design of wear-out sensors for the various aging phenomena (e.g., [25]), they are not currently integrated in commercial devices. For this reason, a commonly-used strategy to estimate the aging status within a system is to employ statistical lifetime reliability models relying on the existence of per-core thermal sensors within the platform (e.g., [18]).

In this work we consider the standard lifetime reliability model for a given system based on a Weibull distribution [26]:

$$R(t) = e^{-\left(\frac{t}{\alpha(T)}\right)^{\beta}} \tag{1}$$

being t the current instant of time (generally measured in hours), T the constant worst-case processor temperature (Kelvin degrees), β the Weibull slope parameter, and $\alpha(T)$ the scale parameter, or aging rate. The $\alpha(T)$ parameter formulation depends on the considered wear-out mechanisms, that are, for instance, the electromigration, the hot carrier injection (HCI), or the thermal cycling. Moreover, the Sum Of Failure Rate (SOFR) approach can be used to combine various failure mechanisms within a single $\alpha(T)$ formula.

Since the temperature of a tile varies in time due to the changes in the resources utilization and the intensity of the executed workload, the Reliability Monitor implements the advanced reliability model defined in [27] supporting temperature changes as follows:

$$R(t) = e^{-\left(\sum_{j=1}^{i} \frac{\tau_j}{\alpha_j(T)}\right)^{\beta}} \tag{2}$$

where τ_j represents the duration of each period of time with constant steady-state temperature T_j starting from time 0 up to time t (i.e., $t = \sum_{j=1}^{i} \tau_i$). From the implementation point of view, the module periodically samples the temperature T_i of each tile by means of the available sensor and computes the corresponding $\alpha_i(T)$. The reliability of the tile is updated by accumulating the duration of the last period τ_i and the computed $\alpha_i(T)$ to the exponent in Equation 2. Finally, all reliability values are collected in a Tile Reliability Matrix transmitted to the Power Controller. Then, a downstream module, called Application-level Reliability Calculator (ARC), aggregates the information contained in the Tile Reliability Matrix on the basis of the Application Matrix to compute the so-called Application Reliability Vector (ARV). Each value in the ARV is computed by multiplying the reliability values of the set of tiles allocated to a single application (i.e., from a reliability point of view, we consider the set allocated tiles as a series system).

As a final note, the overhead caused by the proposed monitoring infrastructure is negligible on both computation and communication. In fact, the monitoring routines only

periodically collect single floating point values and transmit them on the network. Moreover, at each period, the amount of traffic generated for power management equals to approximately 1KB in the considered architecture in Figure 1, which is mainly due to the transmission of the five matrices of floating point values.

V. POWER CONTROLLER UNIT

The novel Power Controller Unit integrated in the overall system is shown at the bottom of Figure 1. The unit is invoked periodically or at specific events, i.e. application's start and termination, and performs a dynamic power management. The mission of this unit is to maximize the system throughput, in terms of executed applications per unit of time and considering possible application QoS requirements, while not violating the available power budget, and, at the same time, to prolong the system lifetime by avoiding unnecessary thermal stress causing and slowing down the aging process in the various processing tiles. To this end, PCPG is used for power gating the cores that are not running any task, while an advanced tuning of per-core DVFS is performed to achieve the specified mission.

Within the Power Controller Unit a Proportional-Integral-Derivative (PID) controller monitors the error between the actual power consumed by the system and the provided TSP over time. Then, the downstream Reliability-aware Power Allocator Unit uses the output of the PID controller together with the data collected from the monitoring units to handle the power management features. In particular, it tunes VF levels of the various tiles with the granularity of the mapped applications to optimize applications' execution times by exploiting available power budget in an efficient way. Moreover, in some specific situations where there is a deficit of power, the unit may also command the RTM unit to kill some running application.

The Reliability-Aware Power Allocator is composed of three different submodules. An Operation Mode Selector monitors the intensity of the current workload experienced by the system and properly selects the corresponding operating modes: over-boosting mode when experiencing an intensive workload or reliability-aware mode in the other cases. Thus, the actual Power Allocator applies specific power allocation policies in the two operating modes: while in the first case, the system is used at full-speed to satisfy the highly demanding user-requests, in the second case power allocation is performed by leveraging also the aging status of processing tiles. Finally, the last submodule, called the Reliability Balancer and invoked only in the reliability-aware mode, performs a minimal refinement of the decisions of the Power Allocator Unit in order to further reduce the stress on the regions of the grid experiencing a fast aging. The overall workflow of the Reliability-Aware Power Allocator is sketched in Algorithm 1, while the internals of the various modules are presented in the following subsections.

A. PID Controller

A classical PID controller is used for controlling the power consumption of the system in a feedback control loop. In particular, the error between the current power consumption

---

Algorithm 1 Reliability-Aware Power Allocator

Inputs: $PID_{out}, RAI, ABUV, AIRV, APV, APUV, ARV, TSP,$
$\qquad newApp\_interrupt, Error, workload$
Output: $V_{PEs}, Freq_{PEs}, killApp$
Global Variables: $DVFSList, I_{set}, NI_{set}, C_{set}, NC_{set}, AppReliability_{set},$
$\qquad PCPowerLimit, timer, curr\_state$
Constant values: $bufferUtilizationLimit, W_{th}, T_{th}$
Body:
1: **if** $curr\_state = reliability\text{-}aware$ and $workload \geq W_{th} + \Delta$ **then**
2: $\quad curr\_state \leftarrow over\text{-}boosting;$
3: **else if** $curr\_state = over\text{-}boosting$ and $workload \leq W_{th} - \Delta$ **then**
4: $\quad curr\_state \leftarrow reliability\text{-}aware;$
5: $(V_{PEs}, Freq_{PEs}, killApp) \leftarrow$
$\qquad powerAllocator(PID_{out}, RAI, ABUV, AIRV, APV,$
$\qquad\qquad APUV, ARV, TSP, newApp\_interrupt, Error);$
6: **if** $curr\_state = reliability\text{-}aware$ **then**
7: $\quad$ **if** $timer \geq T_{th}$ **then**
8: $\quad\quad (V_{PEs}, Freq_{PEs}) \leftarrow reliabilityBalancing(ARV, RAI);$
9: $\quad\quad timer \leftarrow 0;$
10: $\quad timer \leftarrow timer + 1;$

---

and the available TSP is monitored over time to decide the power budget that can be allocated to the current instant of time. The general formula for the PID controller is:

$$PID_{out}(t) = K_p e(t) + K_i \int e(t)\,\mathrm{d}t + K_d \frac{de(t)}{dt} \qquad (3)$$

where $PID_{out}(t)$, $e(t)$, K_p, K_i, and K_d are the controller output, error, proportional gain, integral gain, and derivative gain, respectively. The gains of the PID controller are appropriately adjust after several Matlab simulations. Two key factors were considered in our simulations viz., the system stability and the system robustness against power disturbance. When considering a many-core system featuring runtime application mapping, there are three main events that influence the power trace curve, and therefore $e(t)$: 1) the start of a new application, 2) the end of a running application, and 3) the intra-application variations in the running applications due to the task dependencies and varying switching activities of the various tasks. During our preliminary experimental analysis, we noticed that the PID controller can efficiently support the Power Allocator in making decisions on handling the latter two situations. However, when an application enters the system, a high overshoot may easily happen especially when the application size is large and demands several dark cores to be activated. In fact, the newly incoming application will cause an increase in the instantaneous power consumption. Since this phenomenon causes a drastic drop down in the error $e(t)$ not properly handled by the PID, we introduce a Disturbance Rejector Unit within the Power Allocator to handle such power spikes in a proactive way (See Section 5.3).

B. Operating Mode Selector

The two objectives the Power Controller Unit aims at optimizing, i.e. the system's performance and reliability, are frequently conflicting if not accurately handled. In fact, to achieve higher performance it is necessary to increase VF and consume more power, but this would cause higher temperature hotspots within the chip and a faster aging. In contrast, a reliability-aware power capping approach would unstress the architecture but also worsen the performance. Another interesting difference between the two considered objectives is the highly different time horizon along which they are analyzed. Performance is generally measured in quite short time intervals

Figure 3: State machine diagram of the algorithm within the Operating Mode Selector

Algorithm 2 Power Allocation Algorithm

Inputs: $PID_{out}, RAI, ABUV, AIRV, APV, APUV, TSP,$
$\quad\quad newApp\_interrupt, Error$
Output: $V_{PEs}, Freq_{PEs}, killApp$
Global Variables: $DVFSList, I_{set}, NI_{set}, C_{set}, NC_{set}, appReliability_{set},$
$\quad\quad PCPowerLimit$
Constant values: $bufferUtilizationLimit$
Body:
1: $PCPowerLimit \leftarrow \frac{TSP}{\#activeCores};$
2: $(I_{set}, NI_{set}) \leftarrow IRClassifier(AIRV, RAI);$
3: $(C_{set}, NC_{set}) \leftarrow BUClassifier(ABUV, RAI);$
4: **if** $newApp\_interrupt = true$ **then**
5: $\quad (V_{PEs}, Freq_{PEs}, killApp) \leftarrow$
$\quad\quad\quad proactiveDistRej(error, RAI, ABUV, APV, APUV);$
6: **else**
7: $\quad$ **if** $PID_{out} < 0$ **then**
8: $\quad\quad (V_{PEs}, Freq_{PEs}, killApp) \leftarrow$
$\quad\quad\quad VF_{downscaler}(RAI, ABUV, APV, APUV, PID_{out},$
$\quad\quad\quad\quad PCPowerLimit);$
9: $\quad$ **else**
10: $\quad\quad (V_{PEs}, Freq_{PEs}, killApp) \leftarrow$
$\quad\quad\quad VF_{upscaler}(RAI, ABUV, APV, APUV, PID_{out},$
$\quad\quad\quad\quad PCPowerLimit);$

lasting seconds to few hours. In fact, applications arrive and leave the system with a very fast trend, and consequently management decisions taken on each single application may considerably affect the system's performance. At the opposite, the aging status of a system is a very slow process which is influenced by the average trend in the management decisions over a long period of time.

As discussed in [28], [29], computing systems are generally subject to a variable workload over the time alternating periods of intensive activity and periods with a low number of running applications. Moreover, to fulfill the QoS demands of the running workload it is not always necessary to execute at highest VF levels especially when the system is not experiencing an intensive workload composed of many applications in the ready queue. To face with this variability and handle the two conflicting goals we define two different operating modes, called over-boosting mode, and reliability-aware mode, in which different decision policies are used in the Power Allocator. The *over-boosting mode* is characterized by a highly intensive workload. For this reason, the system needs to work at full speed, and, consequently, reliability issues are ignored. At the opposite, in *reliability-aware mode*, metrics from the reliability monitoring are taken into account to avoid thermal hotspots while trying to obtain good performance as well. Therefore, the ultimate goal of these two operating modes is to obtain in the short term optimal system's performance while mitigating unnecessary stress and wear-out in the architecture in the long term. Therefore, the Power Controller Unit is provided with an Operating Mode Selector which monitors the overall amount of workload the system is experiencing in the current period of time and consequently decides the operating mode to select. The amount of workload is computed by the Operating Mode Selector in percentage value by monitoring the status of the ready queue. Then, the operating mode is transmitted to the Power Allocator that will use the corresponding decision policies. The Operating Mode Selector internally behaves as a Finite State Machine, as shown in Figure 3, that on the basis of a given threshold switch between the over-boosting mode to the reliability-aware one. A tolerance guard band can be used around the threshold value to avoid excessive oscillations between the two different operating modes.

C. Power Allocator

The overall behavior of Power Allocator is represented in Algorithm 2. As the first step, the module pre-processes some input data useful for subsequent elaborations. First, the per-core power limit *PCPowerLimit* is computed by dividing the

TSP in the same portions by the number of active cores in the system (*#activeCores*). When TDP is chosen over TSP, this step can be ignored. Next, running applications are partitioned by means of the *IRClassifier* function in intensive (I_{set}) and non-intensive (NI_{set}) sets in terms of the network intensity, based on *AIRV*. The classification algorithm has been borrowed from [23], where it is possible to find all the implementation details. Similarly, in the second step, applications are partitioned into congested (C_{set}) and non-congested (NC_{set}) sets, based on *ABUV*. In particular, as previously discussed, an application is tagged as congested if the average buffer utilization of its associated routers is larger than a predefined threshold (e.g., 75%). As a conclusion, each application is tagged with a 2-bit label which can get one of these values: *NI_NC* (non-intensive, non-congested), *NI_C* (non-intensive, congested), *I_NC* (intensive, non-congested), and *I_C* (intensive, congested). Applications are classified every time the controller is executed.

After classifications, the event that has triggered the execution of the Power Controller and the current output of the PID controller are analyzed to select a proper DVFS policy. In particular, when Power Controller is woken up by the periodic timer or by an application's completion event, the PID output, PID_{out}, is compared against zero. In case of overshoot (i.e., when $PID_{out} > 0$ that means power consumption exceeding TSP/TDP), $VF_{downscaler}$ function is invoked to scale down VF levels of a subset of selected applications, while in case of undershoot, the $VF_{upscaler}$ function performs VF upscaling. When a $newApp\_interrupt$ is asserted by the runtime mapping unit due to the start of a new application, the Disturbance Rejecter module is invoked to proactively handle the large overshoot caused by the newly running application. The details on the $VF_{downscaler}$ and $VF_{upscaler}$ functions are discussed in the following.

VF$_{downscaler}$: Here, the main idea is to select the candidate applications to be downscaled among the ones with the lowest priority, letting the high priority applications run at a higher QoS level. Algorithm 3 describes the VF downscaling process. The LP_{apps} function is applied on the overall set of running applications ($C_{set} \cup NC_{set}$) to get the subset of applications with the lowest priority, namely $appSet$. The

Algorithm 3 Voltage and Frequency Downscaling Function

Inputs: $RAI, ABUV, APV, APUV, PID_{out}, PCPowerLimit$
Outputs: $V_{PEs}, Freq_{PEs}, killApp$
Variables: $availableApps, targetApp, failedDVFS, appSet$
Body:
1: $availableApps \leftarrow C_{set} \cup NC_{set}$;
2: **while** $availableApps \neq \emptyset$ **do**
3:     $targetApp \leftarrow \emptyset$;
4:     $appSet \leftarrow LP_{apps}(availableApps, RAI)$;
5:     **if** $appSet \cap C_{set} \neq \emptyset$ **then**
6:         $appSet \leftarrow appSet \cap C_{set}$;
7:     **else**
8:         $appSet \leftarrow appSet \cap NC_{set}$;
9:     **if** $current\_state = reliability\text{-}aware$ **then**
10:         $targetApp \leftarrow lowReliability(appSet, appReliability_{set})$;
11:     **else**
12:         $targetApp \leftarrow lowD_{prf-pwr}(appSet, APV, APUV)$;
13:     $(V_{PEs}, Freq_{PEs}, failedDVFS) \leftarrow$
            $DVFS(RAI, targetApp, PID_{out}, PCPowerLimit)$;
14:     **if** $failedDVFS = true$ **then**
15:         $availableApps \leftarrow availableApps \setminus \{targetApp\}$;
16:         **if** $availableApps = \emptyset$ **then**
17:             $killApp \leftarrow targetApp$;
18:             return;
19:     **else**
20:         $DVFSList \leftarrow DVFSList \cup \{targetApp\}$;

Algorithm 4 Voltage and Frequency Upscaling Function

Inputs: $RAI, ABUV, APV, APUV, PID_{out}, PCPowerLimit$
Outputs: $V_{PEs}, Freq_{PEs}, killApp$
Variables: $availableApps, targetApp, failedDVFS, appSet$
Body:
1: $targetApp \leftarrow None$;
2: $availableApps \leftarrow DVFSList$;
3: **while** $targetApp = None$ **do**
4:     $appSet \leftarrow availableApps \cap NC_{set} \cap NI_{set}$;
5:     **if** $appSet = \emptyset$ **then**
6:         $appSet \leftarrow availableApps \cap NC_{set} \cap I_{set}$;
7:         **if** $appSet = \emptyset$ **then**
8:             $appSet \leftarrow availableApps$;
9:     $appSet \leftarrow HP_{apps}(appSet, RAI)$;
10:     **if** $current\_state = reliability\text{-}aware$ **then**
11:         $targetApp \leftarrow highReliablity(appSet, appReliability_{set})$;
12:     **else**
13:         $targetApp \leftarrow highD_{prf-pwr}(appSet, APV, APUV, PID_{out})$;
14:     $(V_{PEs}, Freq_{PEs}, failedDVFS) \leftarrow$
            $DVFS(RAI, targetApp, PID_{out}, PCPowerLimit)$;
15:     **if** $failedDVFS = true$ **then**
16:         $availableApps \leftarrow availableApps \setminus \{targetApp\}$;
17:         $targetApp \leftarrow None$;
18:         continue;
19: $DVFSList \leftarrow DVFSList \setminus \{targetApp\}$;

filtering is performed according to the application priority vector contained in the *RAI*. Then, $appSet$ is even more filtered to consider only the subset of congested applications, with the aim at improving network throughput. In fact, tiles residing in a congested region of the architecture potentially dissipate unnecessary power (particularly static) due to low network throughput. As VF downscaling has also the effect on throttling of packet injection, it can alleviate the network congestion for such applications and save power. In case the selected $appSet$ is empty, congested set is replaced by non-congested set (NC_{set}).

Finally, the target application to be downscaled is selected in the $appSet$ by means of a specific policy of the current operating mode. In case of reliability-aware operating mode, the application characterized by the lowest reliability in the *ARV* is selected in order to reduce the heating in the allocated tiles, consequently unstress the region. While, in case of $over-boosting$ mode, a $lowD_{prf-pwr}$ function is used to select the application with the lowest performance loss to power reduction ratio. The identified target application ($targetApp$) is then downscaled by the $DVFS$ function as per PID_{out} and $PCPowerLimit$.

The downscaling process continues iterating on the $availableApps$ list to look for the next target application. Moreover, the $DVFS$ function will fail when it cannot throttle the target application any further according to the application type. This occurs when voltage and frequency cannot be reduced anymore. When the $failedDVFS$ variable is asserted, the application will be removed from the $availableApps$ and the algorithm continues iterating on the list until an alternative application is found. Finally, if no application can be found to scale down, the last target application to be scaled down will be assigned to $killApp$, so that the RTM unit will kill it.

VF$_{upscaler}$: VF upscaling of processing tiles is presented in Algorithm 4. The algorithm works on the set of already downscaled applications, $DVFSList$, and performs a set of filtering to identify the target application. More precisely, first, $DVFSList$ is filtered to consider only the non-intensive and non-congested applications. In case the result is an empty set,

the filtering is relaxed to consider only the non-intensive tag and, if also this last filter fails obtaining a non-empty set, the original $DVFSList$ content is considered. Finally, the last filtering is performed to select only the applications with the highest priority by means of the HP_{apps} function to consider the QoS demands. The basic idea at the basis of this selection process is that upscaling VF levels of a tile residing in a congested area and having a high injection rate may result in zero performance gain if on-chip communication network is the bottleneck. That is the reason why in VF upscaling process, in contrast with downscaling, a higher priority is given to congestion than application priority in the algorithm.

The target application to be upscaled in selected by means of a specific policy according to the current operating mode. In case of *reliability-aware* operating mode, the application characterized by the highest reliability in the *ARV* is selected. In this way, we will put the stress in the youngest region of the device. While, in case of *over-boosting* mode, a $HighD_{prf-pwr}$ function is used to select the application with the highest performance loss to power increase ratio. The chosen target application is then upscaled by the $DVFS$ function as per PID_{out} and $PCPowerLimit$. The algorithm will continue iterating until an application is not found by removing the attempted candidates from the $availableApps$ list. Finally, the target application is removed from the $DVFSList$.

highD$_{prf-pwr}$ and lowD$_{prf-pwr}$ functions: These functions search for an application with the highest or lowest performance-power ratio, respectively, in a given set to be the target of VF upscaling or downscaling. The original idea of such functions can be found in [11], where product of core utilization ($Util$) and aggregated frequency ($Freq$) is used as a high-level computational capacity metric. In this metric, the frequency is weighted to deduct the idling cycles. We extend this metric by aggregating core utilization in an application ($appUtil$), provided by $APUC$, to calculate the performance of an application as:

$$Perf_{current} = appUtil \times Freq_{current} \quad (4)$$

Then, the performance-power ratio is calculated as:

$$D_{prf-pwr} = \frac{Perf_{next} - Perf_{current}}{Power_{next} - Power_{current}} \quad (5)$$

where $Power_{current}$ is the power consumption of the current application provided by the APC unit, and $Power_{next}$ and $Perf_{next}$ are the estimated power consumption and performance of the application after the DVFS process. The next VF level (V_{dd_next} and $Freq_{next}$) are estimated for the candidate applications based on the magnitude of PID_{out} and application size. $Perf_{next}$ and $Power_{next}$ are calculated as:

$$Perf_{next} = Perf_{current} \times \frac{Freq_{next}}{Freq_{current}} \quad (6)$$

$$Power_{next} = Power_{current} \times \frac{Freq_{next}}{Freq_{current}} \times \left(\frac{V_{dd\_next}}{V_{dd\_current}}\right)^2 \quad (7)$$

After calculating $D_{prf-pwr}$ for all the applications in *appSet*, $lowD_{prf-pwr}$ and $highD_{prf-pwr}$ functions find the application with the lowest and highest $D_{prf-pwr}$ value as the target application for downscaling and upscaling, respectively.

DVFS function: This function actuates the VF downscaling or upscaling of all the tiles allocated for the target application according to the specified PID_{out} and $PCPowerLimit$. In case of downscaling, the minimum throughput specified for the application in the RAI is considered as an additional parameter to identify the minimum VF level applicable (as in [30]). Thus, the function identifies the minimum VF level for the downscaling or the maximum VF level for the upscaling to be applied to all the cores in order to satisfy the three constraints. If no solution can be found, the function returns a failure.

Proactive Disturbance Rejection (PDR): As discussed in the overall description of the approach, the commence of a new application may cause a drastic overshoot in the error between the consumed power and the available TSP/TDP and this phenomenon cannot be properly managed by the PID controller with the functionalities of the Power Allocator discussed so far. For this reason, these sporadic events can be efficiently managed by scaling down a selected set of already running applications in a proactive way in order to collect the power budget required by the new application before its start. The *proactiveDistRej()* function is depicted in Algorithm 5. The strategy first estimates the power consumption of the new application will have, $appPredictedPower$, by using the number of tasks of the application (N) and average power consumed by actively running tiles (P_{avg}), extracted from *RAI* and *APV*. Then, it computes $proactiveError$ as the difference between $Error$ and $appPredictedPower$. Actually, $proactiveError$ represents the PID input without the contribution of the new application supposing that it is already running. Actually, we need such a value since the aim is to collect the necessary power budget by scaling down other applications. Therefore, if $proactiveError$ is positive, there is availability of power budget for the new application, and it is mapped without any further scaling. Conversely, as shown in the second part of the algorithm, if $proactiveError$ is negative, other applications are scaled down. To perform this task, a proportional controller with gain K'_p is used. Here,

---

Algorithm 5 Proactive Disturbance Rejection (*proactiveDistRej()*).

Inputs: $Error, RAI, ABUV, APV, APUV, PCPowerLimit$
Outputs: $V_{PEs}, Freq_{PEs}, killApp$
Variables: $failedDVFS, appPredictedPower, proactiveError, P_{out},$
        P_{avg}
Constant values: K'_p
Body:
1: $P_{avg} \leftarrow computeAvgPowerConsumption(RAI, APV);$
2: $N \leftarrow countTaskOfNewApplication(RAI);$
3: $appPredictedPower \leftarrow N \cdot P_{avg};$
4: $proactiveError \leftarrow Error - appPredictedPower;$
5: **if** $proactiveError < 0$ **then**
6:     $P_{out} \leftarrow K'_p \cdot proactiveError;$
7:     $(V_{PEs}, Freq_{PEs}, killApp) \leftarrow$
        $VF_{downscaler}(RAI, ABUV, APV, APUV, P_{out},$
            $PCPowerLimit);$

---

the integral and derivative terms are removed because when such sporadic rises occur, history-based (i.e., integral term) or prediction-based (i.e., derivative term) decision making will most likely affect the controller's response. Then, the output of the controller (P_{out}) determines the extent by which currently running applications are to be scaled so that the new application can be mapped without violating TSP/TDP. Such an downscaling of the currently running applications is performed as discussed above by using $VF_{downscaler}$ function.

D. Reliability Balancer

Even though when the controller is in *reliability-aware* mode the amount of stress on the cores is used as one of the metrics in scale up/down process, there might be still possibility to further decrease the VF level of the cores with a negligible performance penalty. In this way, we can achieve a more efficient reliability balancing in long term. For this, the last module in the Power Controller is a Reliability Balancer that tries to adjust the distribution of power on chip by marginal changes in VF levels of the tiles to reshape the unbalanced reliability distribution. Algorithm 6 shows reliability balancing function that is running periodically whenever the system operating in *reliability-aware* mode. In the first step, average of all the application reliability values (ARV) is calculated by function $calcAvgRel$. After that, among all the running applications, if the reliability value for one application app_i is lower than the overall average reliability and this application can be downscaled according to its priority and required QoS (by using the $check_{downscaling}$ function), the VF level of the allocated cores will be scaled down by one step. This downscaling provides a power-slack (the difference between current power consumption and maximum upper bound). Hence, after that, one application among $DVFSList$ whose reliability is higher than $avgRel$ is selected to be scaled one step up. The $reliabilityBalancing$ function is managed to be run at coarse interval time (see Line 10 in Algorithm 1) and its time interval is selected to be much longer than RA-$powerAllocator$ to obviate considerable negative effect on the system performance which can be imposed by $reliabilityBalancing$ function. It should be noted that as the size of the applications are different and the relationship between the VF level and power is non-linear, $VF_{onestep\_down}$ scaling down the resources of an application and $VF_{onestep\_up}$ scaling up another application to adjust power might cause

Algorithm 6 Reliability balancing (*reliabilityBalancing()*).

Inputs: ARV, RAI
Outputs: V_{PEs}, $Freq_{PEs}$
Variables: $avgRel$, $usedAppls$
Body:
1: $avgRel \leftarrow calcAvgRel(ARV)$;
2: $usedAppls \leftarrow \emptyset$;
3: **for all** $app_i \in RAI$ **do**
4:    **if** $appRel_i < avgRel$ and $app_i \notin usedAppls$ **then**
5:       **if** $check_{downscaling}(V_{PEs}, Freq_{PEs}, app_i)$ **then**
6:          **for all** $app_j \in DVFSList$ **do**
7:             **if** $app_j \notin usedAppls$ and $i \neq j$ and $appRel_j > avgR$ **then**
8:                $(V_{PEs}, Freq_{PEs}) \leftarrow VF_{onestep\_down}(app_i, RAI)$;
9:                $(V_{PEs}, Freq_{PEs}) \leftarrow VF_{onestep\_up}(app_j, RAI)$;
10:                $usedAppls \leftarrow usedAppls \cup \{app_i, app_j\}$;
11:                break;

---

disturbance in overall power consumption. However, as such disturbance is marginal, it can be handled by Power Allocator in the subsequent iterations.

VI. EXPERIMENTAL EVALUATION

We experimentally evaluated the proposed multi-objective dark silicon aware power management approach by means of an in-house SystemC system-level many-core simulation framework, based on the one used in [12]. The tool is based on Noxim [31], to simulate the network infrastructure. A functional model of the processing tiles is implemented and characterized according to the specifications of the Niagara2 in-order core obtained from McPAT [32]. Physical scaling parameters were extracted from the Lumos framework [33] by considering the 16nm CMOS technology node; in particular, we imported specifications for power modeling, voltage-frequency scaling, and thermal design power (TDP) calculation. In the simulations, we considered a 12×12 many-core having a chip area equal to $138mm^2$. For the DVFS purpose, we use 15 VF levels (similar to Intel SCC) including near-threshold operation extracted from the Lumos framework (the *pessimistic* model). The minimum and maximum VF levels are set to $(0.456V, 300MHz)$ and $(0.908V, 5.2GHz)$, respectively. The frequency of the on-chip communication network (e.g., routers) is set to the maximum level (i.e., 5.2GHz) to demonstrate that even at the maximum NoC speed, the network can get congested and should be taken into account in power management along with the other parameters. For the TSP calculation, we used the chip characterization in [7]. We set ambient temperature to $45°C$, a threshold temperature that triggers thermal management to $80°C$, maximum chip power consumption from the power supply to $300W$, and the power consumption of an inactive core to $0.3W$. Finally, we used Hotspot [34] to calculate the temperature from the power traces at runtime, and, for a proof of concept, for the reliability model we considered the electromigration aging mechanism, characterized as in [27].

Two different application types have been defined: non-realtime, having a low priority, and soft realtime, with a high priority. Several instances of non-realtime applications spanning from 4 to 35 tasks have been generated using TGG [35]; communication and computation volumes are randomly distributed. We model MPEG4 and VOPD multimedia applications as soft realtime applications. We pre-calculate the minimum VF level for soft realtime tasks for their worst-case
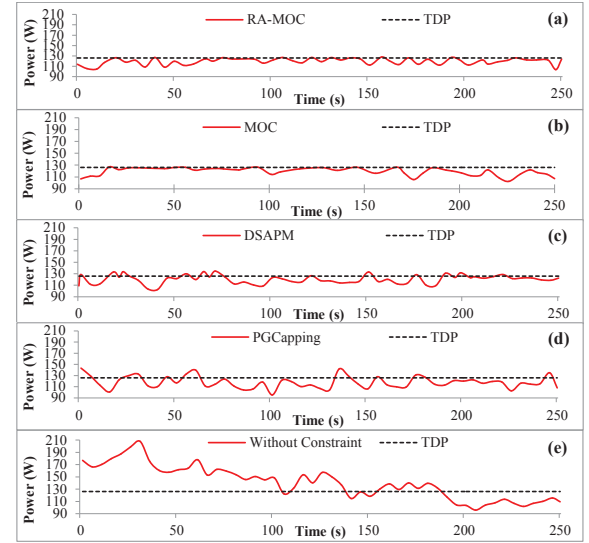

Figure 4: System's power consumption to honor TDP

Figure 5: System's power consumption to honor TSP

contiguous mapping. The workload has been defined in terms of a random sequence of applications entering the incoming application ready queue. The workload mix consists of 30% of soft realtime applications and 70% of non-realtime ones. This sequence is kept fixed in all experiments for the sake of fair comparison.

The runtime resource management layer has been integrated into the described simulator. For the RTM unit, we adopted the state-of-the-art methods SHiC [22] and CoNA [36]. The controller invocation period has been set to $50ms$ as in [11].

In a first experimental campaign, we compare the power management capabilities and performance efficiency of the proposed approach by considering different approaches: 1) our reliability-aware multi-objective controller (*RA-MOC*), 2) our multi-objective controller without reliability consideration (*MOC* as proposed in [12]), 3) *PGCapping* [11] where only core's power-performance ratio is considered as feedback for the PCPG and per-core DVFS actuation, 4) *DSAPM* [17] where no information regarding performance and packet injection rate of tiles is used as feedback, and, finally 5) *without*

TSP/TDP constraint. Without TSP/TDP constraint is the scenario where the system is not limited in terms of maximum power consumption; this is the situation when, in reality, the chip is damaged due to overheating. To perform a fair comparison, we aligned PGCapping and DSAPM techniques to the proposed approach, in particular to use the same 15 VF levels for per-core DVFS. Finally, we consider a 10s warm up phase for the results.

Figure 4 presents the power traces of the system managed by the various considered approaches when honoring a constant TDP set to 126W (the value has been calculated based on the chip power density). Deviation of power consumption from the TDP line reflects either violation or under-utilization of power budget. It is worth mentioning that the power trace is reported after the warm up phase through which the system application entrance rate and power consumption become stable. It is possible to notice that PGCapping, DSAPM and without-constraint power managements mostly tend to overshoot or undershoot from TDP presenting a considerable oscillation, while the proposed approach (with or without reliability management) is able to better exploit the available power budget. Even though PGCapping benefits from the cores' power-performance values, fed back by the controller and thus increases the system throughput to some extent, it suffers from the under-utilization issue as it does not consider the network congestion and applications injection rates. DSAPM considers network congestion, however it also suffers from the under-utilization issue as it is agnostic of cores' performance value and applications' injection rates. Finally, both PGCapping and DSAPM techniques refuse to properly handle occasional overshoots due to new application arrivals. When considering the proposed approach, both MOC and RA-MOC have the best control on the power trace to stay in close proximity with TDP. In cases where power consumption exceeds TDP, the MOC controller rapidly reduces the power consumption by a proper voltage and frequency scaling.

Figure 5 demonstrates the aforementioned power management scenarios to honor dynamic TSP values. As shown in the figure, the conclusions we made for TDP are also valid for dynamic TSP, the MOC-based system is stable even when budget is changed at runtime. It is worth noting that TSP does not radically change (often between 141W and 149W) as the system is mostly busy and the majority of cores are active. To summarize, we computed the percentage of time the power budget is violated by the various power management policies over the overall simulation time. Results are shown in Figure 6. It can be observed that, differently from the state-of-the-art approaches, the proposed MOC approach honors the TDP/TSP constraints for more than 99% of the simulation time.

To assess the performance efficiency of our approach while optimizing the utilization of the power budget, we also analyzed in the same experiment the normalized throughput, in terms of number of completed applications per unit of time, for the considered power management approaches. Results shown in Figure 7 reveal that the proposed MOC method can significantly improve the overall system throughput for different power budget types (up to 29% compared with PGCapping and up to 15% compared with DSAPM). This


Figure 6: TDP/TSP violation for different dynamic power managements (DPM)


Figure 7: Normalized throughput for different dynamic power managements (DPM)

result is obtained due to the advantage of our proposed multi-objective controller which considers both the computation and communication aspects in power management. At the same time RA-MOC has a slightly minor improvement, even if considerably better than state-of-the-art approaches. The motivation is related to the fact that it is also focused on avoiding excessive thermal stress and aging in the architecture thus scarifying a bit on the performance. To show the impact of the power limit on the system performance, we also added the system throughput while no power management technique is applied, (dubbed as *Without Constraint* in the figure); as it can be seen, the dark silicon phenomenon has an impact around 50% on the system performance under the same workload.

In a second experimental campaign, we evaluate the reliability of the tiles during the time. To this purpose, we performed a long-term simulation in which we analyzed the evolution of the lifetime reliability of the various tiles for the overall operational life. To perform an accelerated experiment (with a reasonable simulation time), we enlarged the execution times of the applications to last a few days. Finally, to stimulate the system to switch with a realistic trend between the two different operating modes supported by the proposed approach, we modeled a variable workload during 24 hours, according to the measures on a typical server in Facebook data center from [29]. In order to better show the efficiency of our proposed approach, we considered the proposed approach by using only the reliability-aware power allocator technique without reliability balancing (i.e., without Algorithm 6), namely *RA-MOC-no-balancing*, and the full-fledged approach, namely *RA-MOC*. Using these two steps, we aim at better demonstrating the contribution of reliability-aware power allocation and balancing technique in balancing the overall reliability on the chip. Moreover, we compared RA-MOC and RA-MOC-no-balancing against two state-of-the-art

(a) RA-MOC   (b) RA-MOC-no-balancing   (c) MOC   (d) PGCapping

Figure 8: Effect of reliability-aware power management approach on overall system reliability (TDP-based approach)

(a) RA-MOC   (b) RA-MOC-no-balancing   (c) MOC   (d) PGCapping

Figure 9: Effect of reliability-aware power management approach on overall system reliability (TSP-based approach)

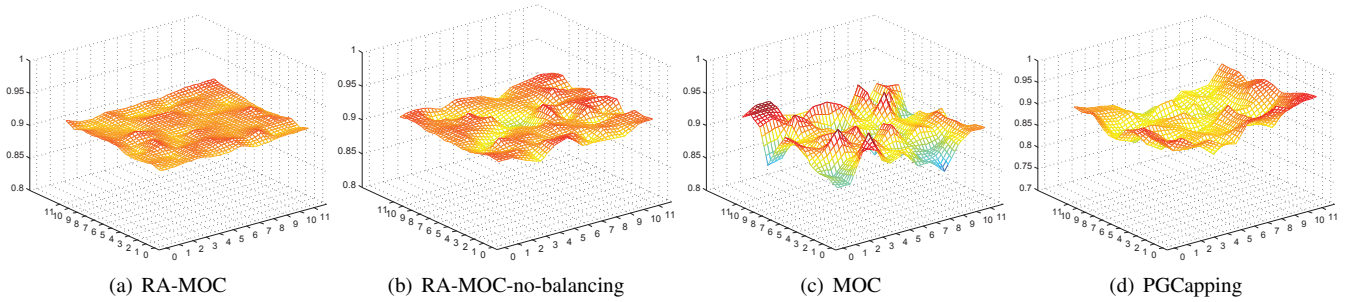(a) RA-MOC   (b) RA-MOC-no-balancing   (c) MOC   (d) PGCapping

Figure 10: Effect of reliability-aware power management approach on cores' reliability after 2 years of system activity (TDP-based approach)

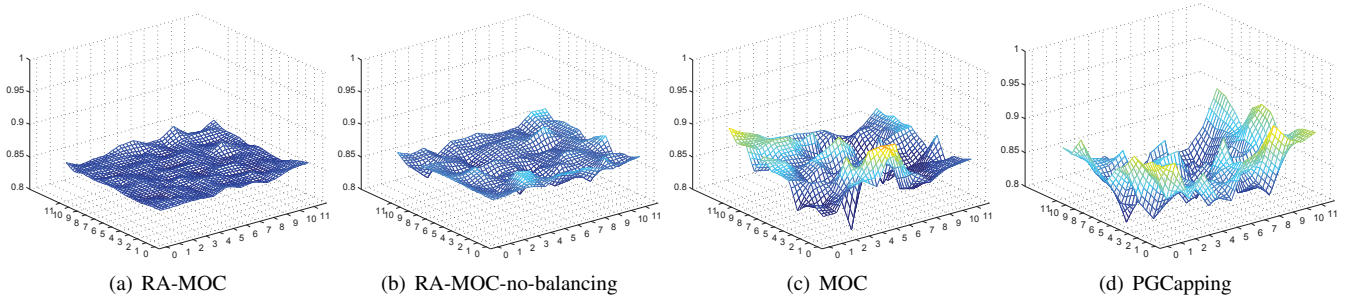(a) RA-MOC   (b) RA-MOC-no-balancing   (c) MOC   (d) PGCapping

Figure 11: Effect of reliability-aware power management approach on cores' reliability after 2 years of system activity (TSP-based approach)

approaches: the original power management strategy without considering reliability [12], called *MOC*, and *PGCapping* [11], which features also lifetime optimization.

Figure 8 and Figure 9 compare the reliability curves of the three approaches while using TDP and TSP as the maximum power upper bound, respectively. Each graph reports the minimum, the maximum, and the average reliability values of the various cores within the architecture over 6 years of activity. When comparing against MOC and PCGating, the proposed approach minimizes the variance in the reliability

values thus maximizing the average lifetime of the various tiles. Instead, since the two state-of-the-art approaches are reliability-agnostic, they distribute the applications without considering the aging values, and therefore lead to an un-balanced distribution of the workload and, consequently, of the aging on the cores. This will lead to a lower reliability of some cores that probabilistically will fail earlier. All these considerations are also confirmed by Table I which reports the Mean-Time-To-Failure (MTTF) of the system computed according to the obtained reliability curves. RA-MOC is able

Table I: System lifetime in terms of MTTF

|	Using TDP	Using TSP
RA-MOC	7.4 years	6.6 years
RA-MOC without RB	6.2 years	6 years
MOC	6 years	5.4 years
PGCapping [11]	6.1 years	5.7 years

Figure 12: Standard deviation of reliability for different dynamic power managements after 2 years system activity

to obtain an improvement in MTTF around 23% and 22% for TDP and TSP-based approaches w.r.t. the reliability-agnostic approach. Finally, if we consider RA-MOC and RA-MOC-no-balancing, it is possible to note that both the two techniques (reliability-aware power allocation and reliability balancing) give a contribution on the the prolonging of the lifetime. RA-MOC-no-balancing presents better results than the reliability-agnostic MOC; then the reliability balancing technique in RA-MOC even more improves such results. From this comparison, it is possible to conclude that the proposed approach is able to balance the reliability of the cores and increase the lifetime with negligible performance penalty comparing to state-of-the-art reliability-agnostic solution (i.e. *MOC*). In fact, MOC selects regions to apply DVFS only by means of performance-centric metrics. In contrast, RA-MOC is able to perform an accurate selection by taking into account also the aging status.

In order to better appreciate the capabilities of the proposed approach, we take a snapshot of the reliability status of the system after 2 years presented in Figure 10 and Figure 11 for TDP and TSP, respectively. As can be seen, the reliability distribution obtained by RA-MOC is more evenly distributed compared to the other scenarios. The worth case is when no reliability consideration is applied. Moreover, it is interesting to notice that cores age faster while using TSP as the power bound compared to using TDP. That is because of the fact that while using TSP, there exists much more utilization compared to while using TDP which results in more power consumption and overall temperature and stresses the cores faster. As a final note, Figure 12 reports the standard deviation regarding the reliability distribution for the aforementioned scenarios in Figure 10 and Figure 11. These data confirm the drawn considerations.

VII. CONCLUSIONS

In this paper, we presented a multi-objective feedback controller approach to manage the power budget among the various processing elements of a many-core system against overshooting of power consumption from a certain limit. The target framework is a NoC-based many-core system using runtime application mapping and experiencing a workload highly variable over time. The feedback to the controller are the processing tiles' power-performance measurements, the tiles aging status, application workloads, and network congestion. Comparing the total system power with the maximum power budget, the controller efficiently changes voltage and frequency of appropriate tiles to optimize performance while prolonging the lifetime of the system by avoiding stress and thermal hotspots. The results showed enhanced system's throughput, TDP/TSP violation and overall lifetime for the proposed platform when compared to state-of-the-art power management policies. Future directions will consider the integration of more advanced power management policies acting also on the communication sub-system and reliability-aware mapping strategies to better prolong lifetime while not incurring much performance penalty. Finally, the proposed framework will be also deployed and validated on real many-core platforms.

REFERENCES

[1] N. Goulding-Hotta, J. Sampson, G. Venkatesh, S. Garcia, J. Auricchio, P. Huang, M. Arora, S. Nath, V. Bhatt, J. Babb, S. Swanson, and M. Taylor, "The GreenDroid Mobile Application Processor: An Architecture for Silicon's Dark Future," *IEEE Micro*, vol. 31, no. 2, pp. 86–95, 2011.

[2] H. Esmaeilzadeh, E. Blem, R. St. Amant, K. Sankaralingam, and D. Burger, "Dark Silicon and the End of Multicore Scaling," *IEEE Micro*, vol. 32, no. 3, pp. 122–134, May 2012.

[3] A. Rahmani, P. Liljeberg, A. Hemani, A. Jantsch, and H. Tenhunen, *The Dark Side of Silicon*, 1st ed. Springer, 2016.

[4] S. I. Association *et al.*, "Int. technology roadmap for semiconductors (ITRS), 2011 edition," 2011.

[5] Y. Xiang, T. Chantem, R. Dick, X. Hu, and L. Shang, "System-level reliability modeling for MPSoCs," in *Conf. Hardware/Software Codesign and System Synthesis (CODES)*, 2010, pp. 297–306.

[6] L. Wang and K. Skadron, "Implications of the Power Wall: Dim Cores and Reconfigurable Logic," *IEEE Micro*, vol. 33, no. 5, pp. 40–48, 2013.

[7] S. Pagani, H. Khdr, W. Munawar, J.-J. Chen, M. Shafique, M. Li, and J. Henkel, "TSP: Thermal Safe Power: Efficient Power Budgeting for Many-core Systems in Dark Silicon," in *Proc. Int. Conf. on Hardware/Software Codesign and System Synthesis (CODES)*, 2014, pp. 10:1–10:10.

[8] P. Bogdan, R. Marculescu, and S. Jain, "Dynamic Power Management for Multidomain System-on-chip Platforms: An Optimal Control Approach," *ACM Trans. Design Automation for Electronic Systems*, vol. 18, no. 4, pp. 46:1–46:20, 2013.

[9] R. David, P. Bogdan, R. Marculescu, and U. Ogras, "Dynamic Power Management of Voltage-Frequency Island Partitioned Networks-on-Chip Using Intel Sing-Chip Cloud Computer," in *Proc. Int. Symp. on Networks-on-Chip (NOCS)*, 2011, pp. 257–258.

[10] T. Muthukaruppan, M. Pricopi, V. Venkataramani, T. Mitra, and S. Vishin, "Hierarchical power management for asymmetric multi-core in dark silicon era," in *Proc. Design Automation Conf. (DAC)*, 2013, pp. 1–9.

[11] K. Ma and X. Wang, "PGCapping: Exploiting Power Gating for Power Capping and Core Lifetime Balancing in CMPs," in *Proc. Int. Conf. on Parallel Architectures and Compilation Techniques (PACT)*, 2012, pp. 13–22.

[12] A.-M. Rahmani, M.-H. Haghbayan, A. Kanduri, A. Weldezion, P. Liljeberg, J. Plosila, A. Jantsch, and H. Tenhunen, "Dynamic Power Management for Many-Core Platforms in the Dark Silicon Era: A Multi-Objective Control Approach," in *Proc. Int. Symp. on Low Power Electronics and Design (ISLPED)*, 2015, pp. 1–6.

[13] M. Khan, M. Shafique, and J. Henkel, "Power-efficient accelerator allocation in adaptive dark silicon many-core systems," in *Proc. Int. Conf. on Design, Automation & Test in Europe (DATE)*, 2015, pp. 916–919.

[14] ——, "Hierarchical power budgeting for Dark Silicon chips," in *Proc. Int. Symp. on Low Power Electronics and Design (ISLPED)*, 2015, pp. 213–218.

[15] M. Shafique, B. Vogel, and J. Henkel, "Self-adaptive hybrid Dynamic Power Management for many-core systems," in *Proc. Int. Conf. on Design, Automation & Test in Europe (DATE)*, 2013, pp. 51–56.

[16] Z. Chen and D. Marculescu, "Distributed Reinforcement Learning for Power Limited Many-core System Performance Optimization," in *Proc. Int. Conf. on Design, Automation & Test in Europe (DATE)*, 2015, pp. 1521–1526.

[17] M.-H. Haghbayan, A.-M. Rahmani, A. Weldezion, P. Liljeberg, J. Plosila, A. Jantsch, and H. Tenhunen, "Dark silicon aware power management for manycore systems under dynamic workloads," in *Proc. Int. Conf. on Computer Design (ICCD)*, 2014, pp. 509–512.

[18] P. Mercati, F. Paterna, A. Bartolini, L. Benini, and T. Rosing, "Dynamic variability management in mobile multicore processors under lifetime constraints," in *Proc. Int. Conf. on Computer Design (ICCD)*, 2014, pp. 448–455.

[19] T. Kim, B. Zheng, H.-B. Chen, Q. Zhu, V. Sukharev, and S. X.-D. Tan, "Lifetime Optimization for Real-time Embedded Systems Considering Electromigration Effects," in *Proc. Int. Conf. on Computer-Aided Design (ICCAD)*, 2014, pp. 434–439.

[20] J. Howard *et al.*, "A 48-Core IA-32 message-passing processor with DVFS in 45nm CMOS," in *Proc. Int. Solid-State Circuits Conference Digest of Technical Papers (ISSCC)*, 2010, pp. 108–109.

[21] Kalray, "Kalray MPPA Manycore," http://www.kalrayinc.com/.

[22] M. Fattah, M. Daneshtalab, P. Liljeberg, and J. Plosila, "Smart hill climbing for agile dynamic mapping in many-core systems," in *Proc. Design Automation Conf. (DAC)*, 2013, pp. 1–6.

[23] K. Chang, R. Ausavarungnirun, C. Fallin, and O. Mutlu, "HAT: Heterogeneous Adaptive Throttling for On-Chip Networks," in *Proc. Int. Symp. on Computer Architecture and High Performance Computing (SBAC-PAD)*, 2012, pp. 9–18.

[24] A. Weldezion, M. Grange, D. Pamunuwa, Z. Lu, A. Jantsch, R. Weerasekera, and H. Tenhunen, "Scalability of network-on-chip communication architecture for 3-D meshes," in *Proc. Int. Symp. on Networks-on-Chip (NOCS)*, 2009, pp. 114–123.

[25] J. Blome, S. Feng, S. Gupta, and S. Mahlke, "Self-calibrating online wearout detection," in *Proc. Int. Symp. on Microarchitecture (MICRO)*, 2007, pp. 109–122.

[26] JEDEC Solid State Tech. Association, "Failure mechanisms and models for semiconductor devices," 2010, JEP122G.

[27] C. Bolchini, M. Carminati, M. Gribaudo, and A. Miele, "A lightweight and open-source framework for the lifetime estimation of multicore systems," in *Proc. Int. Conf. on Computer Design (ICCD)*, 2014, pp. 166–172.

[28] L. Barroso and U. Holzle, "The Case for Energy-Proportional Computing," *Computer*, vol. 40, no. 12, pp. 33–37, 2007.

[29] O. Bilgir, M. Martonosi, and Q. Wu, "Exploring the Potential of CMP Core Count Management on Data Center Energy Savings," in *Workshop on Energy Efficient Design*, 2011.

[30] C. Silvano, G. Palermo, S. Xydis, and I. Stamelakos, "Voltage Island Management in Near Threshold Manycore Architectures to Mitigate Dark Silicon," in *Proc. Int. Conf. on Design, Automation & Test in Europe (DATE)*, 2014, pp. 1–6.

[31] V. Catania, A. Mineo, S. Monteleone, M. Palesi, and D. Patti, "Noxim: An open, extensible and cycle-accurate network on chip simulator," in *Proc. Int. Conf. on Application-specific Systems, Architectures and Processors (ASAP)*, 2015, pp. 162–163.

[32] S. Li, J. H. Ahn, R. Strong, J. Brockman, D. Tullsen, and N. Jouppi, "McPAT: An integrated power, area, and timing modeling framework for multicore and manycore architectures," in *Proc. Int. Symp. on Microarchitecture (MICRO)*, 2009, pp. 469–480.

[33] L. Wang and K. Skadron, "Dark vs. Dim Silicon and Near-Threshold Computing Extended Results," in *University of Virginia Department of Computer Science Technical Report TR-2013-01*, 2012.

[34] K. Skadron, M. Stan, K. Sankaranarayanan, W. Huang, S. Velusamy, and D. Tarjan, "Temperature-aware microarchitecture: Modeling and implementation," *ACM Trans. Archit. Code Optim.*, pp. 94–125, Mar 2004.

[35] "TGG: Task Graph Generator," http://sourceforge.net/projects/taskgraphgen/, last Update: 2013-04-11.

[36] M. Fattah, M. Ramirez, M. Daneshtalab, P. Liljeberg, and J. Plosila, "CoNA: Dynamic application mapping for congestion reduction in many-core systems," in *Proc. Int. Conf. on Computer Design (ICCD)*, 2012, pp. 364–370.

Amir M. Rahmani received his M.Sc. from University of Tehran, Iran, in 2009 and Ph.D. degree from Department of IT, University of Turku, Finland, in 2012. He also received his MBA jointly from Turku School of Economics and European Institute of Innovation & Technology (EIT) ICT Labs, in 2014. He is currently a University Teacher at the University of Turku, Finland. His research interests span energy-efficient and dependable computing, parallel and distributed systems, and Internet-of-Things. He is currently co-leading three Academy of Finland projects entitled "MANAGE", "InterSys", and "SPA".

M. Hashem Haghbayan received the BA degree in computer engineering from Ferdowsi University of Mashhad and the MS degree in computer architecture from University of Tehran, Iran. Since 2014 he is PhD student in University of Turku, Finland. His research interests include high-performance energy-efficient architectures, power management techniques, and online/offline testing. He has several years of experience working in industry as well as developing research tools before starting his PhD.

Antonio Miele is an Assistant Professor at Politecnico di Milano since 2014. He holds a MS degree in Computer Engineering (2006) and a Ph.D. in Information Technology (2010) from Politecnico di Milano, and a M.Sc. in Computer Science (2006) from the University of Illinois at Chicago. His research interests include design and analysis of dependable embedded systems, focusing on HW/SW hardening techniques, reliable adaptive systems, design space exploration w.r.t. reliability, and fault simulation.

Pasi Liljeberg received the MSc and PhD degrees in electronics and information technology from the University of Turku, Turku, Finland, in 1999 and 2005, respectively. He is an adjunct professor in embedded computing architectures at the University of Turku, Embedded Computer Systems laboratory. During the period 2007-2009, he held an Academy of Finland researcher position. Liljeberg is the leader of the Internet-of-Things for Healthcare (IoT4Health) research group.

Axel Jantsch received the Dipl.Ing. and Dr.Tech. degrees from the Technical University of Vienna, Vienna, Austria. He was a professor of electronic system design with the Royal Institute of Technology, Stockholm, Sweden. He is currently a professor in system on chip with TU Wien, Vienna, Austria. His current research interests include VLSI design and synthesis, system-level specification, modeling and validation, HW/SW co-design and co-syntheses, reconfigurable computing, and networks-on-chip.

Hannu Tenhunen received the diplomas from the Helsinki University of Technology, Finland, 1982, and the PhD degree from Cornell University, NY, 1986. In 1985, he joined the Signal Processing Laboratory, Tampere University of Technology, Finland, as an associate professor and later served as a professor and department director. Since 1992, ha has been a professor at the KTH Royal Institute of Technology, Sweden, where he also served as a dean. He has more than 600 publications and 16 patents.