

Alma Mater Studiorum Università di Bologna
Archivio istituzionale della ricerca

Modeling and Controlling Many-Core HPC Processors: an Alternative to PID and Moving Average Algorithms

This is the final peer-reviewed author's accepted manuscript (postprint) of the following publication:

Published Version:

Bambini, G., Ottaviano, A., Conficoni, C., Tilli, A., Benini, L., Bartolini, A. (2026). Modeling and Controlling Many-Core HPC Processors: an Alternative to PID and Moving Average Algorithms. ACM TRANSACTIONS ON AUTONOMOUS AND ADAPTIVE SYSTEMS, 21(2), 1-32 [10.1145/3694687].

Availability:

This version is available at: <https://hdl.handle.net/11585/1013611> since: 2025-04-01

Published:

DOI: <http://doi.org/10.1145/3694687>

Terms of use:

Some rights reserved. The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>).
When citing, please refer to the published version.

(Article begins on next page)

This is the accepted manuscript of:

Giovanni Bambini, Alessandro Ottaviano, Christian Conficoni, Andrea Tilli, Luca Benini, and Andrea Bartolini. 2024. Modeling and Controlling Many-Core HPC Processors: an Alternative to PID and Moving Average Algorithms. ACM Trans. Auton. Adapt. Syst. Just Accepted (September 2024). <https://dl.acm.org/doi/10.1145/3694687>

©ACM 2024. This version of the work is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in <https://doi.org/10.1145/3694687>

Modeling and Controlling Many-Core HPC Processors: an Alternative to PID and Moving Average Algorithms

GIOVANNI BAMBINI, University of Bologna, Italy

ALESSANDRO OTTAVIANO, ETH Zurich, Switzerland

CHRISTIAN CONFICONI, University of Bologna, Italy

ANDREA TILLI, University of Bologna, Italy

LUCA BENINI, ETH Zurich, Switzerland and University of Bologna, Italy

ANDREA BARTOLINI, University of Bologna, Italy

The race towards performance increase and computing power has led to chips with heterogeneous and complex designs, integrating an ever-growing number of cores on the same monolithic chip or chiplet silicon die. Higher integration density, compounded with the slowdown of technology-driven power reduction, implies that power and thermal management become increasingly relevant. Unfortunately, existing research lacks a detailed analysis and modeling of thermal, power, and electrical coupling effects and how they have to be jointly considered to perform dynamic control of complex and heterogeneous Multi-Processor System on Chips (MPSoCs). To close the gap, in this work, we first provide a detailed thermal and power model targeting a modern High Performance Computing (HPC) MPSoC. We consider real-world coupling effects such as actuators' non-idealities and the exponential relation between the dissipated power, the temperature state, and the voltage level in a single processing element. We analyze how these factors affect the control algorithm behavior and the type of challenges that they pose. Based on the analysis, we propose a thermal capping strategy inspired by Fuzzy control theory to replace the state-of-the-art PID controller, as well as a root-finding iterative method to optimally choose the shared voltage value among cores grouped in the same voltage domain. We evaluate the proposed controller with model-in-the-loop and hardware-in-the-loop co-simulations. We show an improvement over state-of-the-art methods of up to 5× the maximum exceeded temperature while providing an average of 3.56% faster application execution runtime across all the evaluation scenarios.

CCS Concepts: • **Hardware** → **Temperature control**; **Chip-level power issues**; **Process, voltage and temperature variations**; **Temperature optimization**; • **Computer systems organization** → *System on a chip*; Real-time system architecture; • **Computing methodologies**;

Additional Key Words and Phrases: Modeling, Control, Nonlinear systems

ACM Reference Format:

Giovanni Bambini, Alessandro Ottaviano, Christian Conficoni, Andrea Tilli, Luca Benini, and Andrea Bartolini. 2024. Modeling and Controlling Many-Core HPC Processors: an Alternative to PID and Moving Average Algorithms . *ACM Trans. Autonom. Adapt. Syst.* 1, 1, Article 1 (January 2024), 31 pages. <https://doi.org/10.1145/>

Authors' addresses: Giovanni Bambini, University of Bologna, Viale del Risorgimento 2, Bologna, Italy, giovanni.bamibini2@unibo.it; Alessandro Ottaviano, ETH Zurich, Gloriastrasse 85, Zurich, Switzerland, aottaviano@ethz.ch; Christian Conficoni, University of Bologna, Viale del Risorgimento 2, Bologna, Italy, christian.conficoni3@unibo.it; Andrea Tilli, University of Bologna, Viale del Risorgimento 2, Bologna, Italy, christian.conficoni@unibo.it; Luca Benini, ETH Zurich, Gloriastrasse 85, Zurich, Switzerland and University of Bologna, Viale del Risorgimento 2, Bologna, Italy, lbenini@iis.ee.ethz.ch; Andrea Bartolini, University of Bologna, Viale del Risorgimento 2, Bologna, Italy, a.bartolini@unibo.it.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

1556-4665/2024/1-ART1 \$15.00
<https://doi.org/10.1145/3694687>

3694687

1 INTRODUCTION

Modern high-performance Central Processing Units (CPUs) are heterogeneous architectures that feature many processing elements integrated on a monolithic silicon die or silicon chiplet. These complex and heterogeneous architectures require a runtime dynamic thermal and power management strategy that allows fine-grained monitoring and control of the application workloads to either keep the system within well-specified Thermal Design Power (TDP) and minimize power spikes and temperature hot spots. Growing attention is indeed being paid to energy-efficient large-scale HPCs.

The control problem introduced by modern many-core architecture is challenging: it is a non-linear problem with fast dynamics, involving multiple coupled control elements that are subjected to variable constraints within a system with high-amplitude and high-frequency disturbances (Section 4.2). These constraints have to be enforced while aiming at selecting the operating point at the highest power-affordable performance for the system executing an application workload. Furthermore, controlled inputs may be coupled due to the many-cores design that pushes the limit of hardware integration for voltage regulators, forcing several cores to share a single voltage through the same power domain rail.

Traditional dynamic power management strategies range from architectural techniques involving gate-level clock/power gating to *runtime active control* carried out by higher layers of the Operating System (OS) and firmware software stacks [56]. In the latter scenario, which is the focus of this work, the control signal travels through several control blocks in cascade steps, as shown in Figure 1. Each control block is characterized by a certain level of system abstraction and knowledge of the surroundings. In the current literature (Section 2), we identify three main categories of controllers and algorithms according to their execution context: a class of controllers that oversee multiple systems (e.g. whole data centers, or entire nodes), a class where the control action originates within a single system through high-level controllers, such as the Operating System Power Management (OSPM) running on the application-class cores or on the off-chip Board Management Controller (BMC), and a class where the control action originates from a dedicated unit integrated within the MPSoCs. The latter are typically low-power embedded microcontrollers in the MPSoC's *uncore* domain, with privileged access and high bandwidth control over the system architecture. Hence, they are referred to as Low-Level Controller (LLC) in this work [32, 33], while the other two categories are referred to as global and local High-Level Controllers (HLCs).

Most of the State-of-the-Art (SotA) focuses on HLCs [8, 19]. We conjecture this happens because of the lack of an open-source hardware (HW)/software (SW) research framework able to capture the complexity and interfaces of an LLC subsystem integrated into a modern HPC MPSoC.

LLCs have also become increasingly important in the context of advanced power management. It has become clear the trend, originating from industry, of abstracting power and system management tasks away from the HLCs Application-class Processors (APs) towards the LLCs [4, 52] intended as control systems that are closer to the controlled hardware components, and increasingly performant LLCs, such as parallel microcontrollers, are able to address multi-variable optimization scenarios without the need to reduce the complexity of the control problem, simultaneously retaining low-power, and energy-efficiency and real-time bounds required by the *uncore* domain and the nature of the control problem [17]. The objective of this work is to provide an overview and analysis of the LLC control problem and its challenges, as well as the design of an open-source algorithm developed in a real-world setting to serve as proxy for future research.

To the best of the authors' knowledge, little research has addressed the control problem of HPC systems with full coupling between power, frequency, voltage, and temperature. As shown in




Fig. 1. Architecture of the cascade control blocks in an HPC system. The three blocks have separate time domains, their own abstraction of the system to be controlled, and distinct external information depending on their scope. In particular, the LLC (inner light green block) controls the low-level, physical parts of the system (Voltage Regulator Module (VRM), Low-Dropout Regulator (LDO), Dispatch Throttling (DT), and Phase Locked Loop (PLL)) based on the target operating points given by the local HLC. The LLC also fetches information from the Process, Voltage, Temperature (PVT) and other sensors and communicates it to the HLC. The HLC (middle dark green block) takes information from the LLC about the system as well as information from the executing application and computes a set of operating points. Finally, the global HLC (external light blue block) takes information from the system and from the surroundings and communicates it to a performance/energy efficiency target to each local HLC.

Section 2, SotA controllers from industry and academia employ simplified control policies often unable to address the challenges of novel heterogeneous many-core architectures.

In this manuscript, we significantly extend the conference work in Bambini et al. [7] and propose the first analysis and implementation of a cascade reactive control policy that tackles the whole multi-variable control problem with actuators coupling and non-idealities. We provide a cost assessment based on Model in the Loop (MIL) and Hardware in the Loop (HIL) emulations. In the latter, we implement the algorithm on a Field-Programmable Gate Array (FPGA)-System on Chip (SoC) framework codesigned with an open-source, RISC-V embedded LLC [39] to account for the physical state time scale of the control variables. This work provides the following contributions:

- We comprehensively describe a multi-variable power and thermal control problem in a modern HPC MPSoC (Section 4 and Section 5). The control targets a LLC that reacts to high-level decisions from the surrounding HLCs (OSPM, BMC) and dispatches Dynamic Voltage and Frequency Scaling (DVFS) operating points while a workload is being executed on the controlled system. We present how to model the computational part of that system considering all the couplings occurring within the same component and between neighboring components.
- We extend the control algorithm presented in [6, 7] to account for variable coupling and introduce a thermal capping approach inspired by the Fuzzy control theory. Finally, we address the shared voltage coupling with an iterative, root-finding method. The resulting improved reactive control policy outperforms the SotA reactive PID under multiple different scenarios (Section 7) in a MIL-based emulation framework.

- We evaluate the feasibility of the designed algorithms on a RISC-V open-source LLC controller (Section 7). The HW controller is mapped on a modern Heterogeneous System on Chip (HeSoC) FPGA and relies on a HIL evaluation framework that includes the controller and a power/thermal model of the controlled plant.

This work is freely available at the following repositories:

- https://github.com/Ev3nt1ne/AechPeSi_lab: the Matlab simulation environment and control part;
- <https://github.com/pulp-platform/control-pulp> and https://github.com/pulp-platform/control_pulp_pcf: the proposed parallel LLC controller, including open-source hardware, firmware, and control algorithm;
- https://github.com/pulp-platform/pulp_hpc_cosim: the simulation environment written in C for FPGA-based co-simulation.

2 RELATED WORKS

The joint control objective of HLC and LLC controllers is to perform power and thermal management on the system while maintaining a dynamic trade-off between maximum performance and energy efficiency [50]. Nevertheless, they differ in their computing capabilities, safety requirements, and the accessibility to sensors and actuators surrounding the physical system under control [60].

The HLC is generally part of the OS routine, or a user-level application executing concurrently with the main system workload [10]. Leveraging directly the application-class Processing Elements (PEs), the HLC has access to higher computing capabilities than the LLC [39], albeit limited by the contention with the executing workload. Furthermore, lower processor privilege levels (e.g., supervisor in the RISC-V community [45]) have typically restricted access to the system itself for security reasons.

For these reasons, HLCs generally support complex control algorithms aimed at *prediction* and *optimality*, which take into account a comprehensive characterization of the surroundings of the controlled system, such as coupling with neighboring systems, the ambient temperature, and other external factors [19], instead of focusing on specific aspects of the system itself. Their objective is shifted more towards power efficiency and optimal application performance trade-off than power and thermal capping. To this end, HLCs use an abstraction of the physical system through common protocols such as Advanced Configuration and Power Interface (ACPI) [52].

Model Predictive Control (MPC) has been used to control the temperature, by computing an optimal trajectory based on well-crafted energy cost functions [8, 34, 56]. Model identification techniques [12, 22] including ones based on Machine Learning (ML) [35, 44] are largely used to develop complex and accurate HLC controllers that are updated in real-time. A wide literature exists for ML-based HLC controllers which not only optimize energy efficiency and execution performance but also task allocation and threads management [41].

A trend toward lightweight HLC controllers exists, but it is mainly related to mobile platforms which differ significantly from HPC systems by number of PEs and power scales. In Bhat et al. [13], a lightweight Dynamic Power and Thermal Management (DPTM) based on an offline system model identification is deployed on a mobile platform to improve performance and energy efficiency by leveraging the sensitivity of the workload to frequency changes. Fuzzy controllers have also been successfully used in this context [18, 36].

While advanced research in this field focuses on HLC, they traditionally rely on LLCs to apply their policies and to carry out fine-grain thermal and power capping [3]. As a matter of fact, end-to-end analysis and evaluations of LLCs are scarce in SotA literature due mostly to the work being done behind industry secret, and the lack of freely accessible LLC hardware, software, and

emulation frameworks. LLCs are indeed tailored to the system they are integrated into — hence, more difficult to generalize — but control abstraction and analysis are possible. This work focuses only on the analysis of LLCs in the HPC application domain.

To the best of the authors' knowledge, there are four major HPC actors in the industry market employing their own LLC: Intel, AMD, IBM, and ARM.

Intel uses individual functional blocks' power consumption as the controlled variable and performs dynamic budget allocation to the various components [21, 57]. The power is computed through a model whose dependencies include leakage information, voltage, temperature, and an estimated workload. The control is characterized by an Exponential Weight Moving Average (EWMA) algorithm over multiple time windows of the order of seconds. According to [26], other thermal control features act in parallel with the power capping. One of Intel's flagship management techniques is Running Average Power Limit (RAPL), a hardware feature for fine-grained energy consumption monitoring and power/thermal capping with support for several power domains in the system, covering CPUs, Graphic Processing Units (GPUs), and Dynamic Random Access Memory (DRAM) controllers. Monitoring happens through model-specific registers (MSRs), simple counters integrated within the architecture.

According to [1, 14], AMD relies on a distributed Proportional Integral Derivative (PID) control structure. Both the thermal and the power control are performed independently using the frequency as the controlled variable, while a voting-box mechanism is tasked to resolve the coupling among components. Given the similarity to Intel architectures, AMD's power and thermal management strategy resembles RAPL from an implementation point of view.

IBM releases its Power9 OpenPOWER firmware as open-source [29]. From the source code, a voting-box mechanism selects the most constraining frequency for the system, similar to the AMD strategy. The temperature and power consumption are controlled independently, with the frequency being the controlled variable. Differently from AMD, IBM's control firmware has a centralized design, i.e. it gathers data to a central unit and performs one single algorithm computation using moving average filters and PIDs. One of the blocks entering the voting-box mechanism is based on workload estimation.

Likewise, Arm releases the System Control Processor (SCP) firmware as open-source, which consists of a cascade design with weight-based power distribution. The thermal dispatching layer precedes the power reduction one. They use power as the controlled variable, a PID for thermal capping, and a two-stage power distribution with a power accumulation variable to distribute unused power to the core.

Our original control algorithm [6], [7] is based on a cascade structure and it uses the cores' estimated power consumption as the controlled variable. Contrarily to ARM, the power budget is firstly distributed among cores with a heuristic algorithm based on the temperature state of each core, then this granted power is further reduced to meet the thermal set-point through a distributed PID layer. Albeit being employed in fully-fledged modern CPUs, such algorithms rely on control strategies that do not account for the general problem of control variable coupling. We believe that an analysis of this scenario is required in the research community, provided the recent availability of energy-efficient yet powerful parallel embedded microcontrollers that can sustain more complex control strategies.

A summary of the analyzed literature is presented in Table 1.

3 BACKGROUND

As fundamental units of modern supercomputers, HPC processors are nowadays heterogeneous systems comprising several computing chiplets connected with an *interposer* by means of through silicon vias (TSVs) [37] and integrating on-chip HBM memories.

Table 1. Summary of existing HLC and LLC SotA implementation in industry and academia. The $\dagger$ symbol in the Step Time column indicates the presence of event-based controller features. VBA stands for Voting Box Algorithm.

Reference	Year	Type	Target	Controller type	Step Time in ms
Academia					
Bartolini et al. [8]	2013	HLC	HPC	MPC	1-10
Bhat et al. [13]	2017	HLC	Mobile	Model-Based	100 $\dagger$
Cui et al. [18]	2017	HLC	Mobile	Fuzzy	1 – 20
Majumdar et al. [34]	2017	HLC	Desktop / Mobile	MPC	$\dagger$
Rapp et al. [44]	2017	Boost	HPC	ML	1
Leva et al. [31]	2018	HLC	Desktop	PI	$\dagger$
Moazzemi et al. [36]	2019	HLC	Mobile	Fuzzy	200
Mandal et al. [35]	2020	HLC	Mobile	ML	50
Tilli et al. [56]	2022	HLC	HPC / Desktop	MPC	1
Bambini et al. [6]	2022	LLC	HPC	PI + Heuristic	0.5
This work	2024	LLC	HPC	Hybrid	0.5
Industry					
PCU (Intel)	in use	LLC	HPC / Desktop / Mobile	PI VBA	0.5
SCP, MCP (ARM)	in use	LLC	HPC / Desktop / Mobile	PI VBA	n.d.
SMU (AMD)	in use	LLC	HPC / Desktop / Mobile	PI VBA	n.d.
OCC (IBM)	in use	LLC	HPC	PI VBA	0.25

In the realm of power and thermal management, a precise model capturing its physical dynamics is of paramount importance. In this work, we focus on modeling a computing chiplet integrating the main application class PEs as the most significant sources of power and temperature fluctuations. From a control perspective, each PE can be considered as a Multiple-Input Multiple-Output (MIMO) system where the supply voltage V_i and the operation frequency F_i are the controlled *input* variables, while the temperature T_i and the power consumption P_i are the *output* variables, as shown in Figure 2. Frequency and Voltage are linked by a functional inequality constraint, as described more in detail in Section 3.2, but they retain a degree of independence. These four variables are the main elements of the control problem.

The power consumption $P_i(t)$ of each PE component c_i of the system can be described by an algebraic non-linear function that links the frequency $F_i(t)$, the voltage $V_i(t)$, and the temperature $T_i(t)$ with the workload $\omega_i(t)$ (i.e. the application) executing on the specific component itself. Different workloads activate different gates of the PE, producing different power outputs. Frequency and workload scale linearly, voltage and temperature have a non-linear relationship [28, 54]. Differently from power consumption, the temperature T_i of each component is a state of the system evolving in time subjected to $P_i(t)$ and nearby temperatures of neighboring PEs [8, 56].

Figure 3 shows the architecture we consider for deriving the mathematical description of the heat dissipation dynamic system. It represents a chiplet integrated on a silicon die over a carrier Printed Circuit Board (PCB), with a copper heat spreader placed over the active silicon devices to




Fig. 2. Representation of a PE-component with the four main elements of the control problem (the inputs Frequency (F) and Voltage (V), and the outputs Temperature (T) and consumed Power (P)) and the workload $\omega_i(t)$. The power consumption P^* is grayed out because it is not directly measurable per single PE. Instead, the power consumption measure is provided for each power rail (i.e. group of PEs).




Fig. 3. Representation of the HPC Processor thermal architecture. The die contains all the PEs which are the source of the heat production through power dissipation. The main heat dissipation path is indicated by the red arrow passing through the Heat Spreader, the Heat sink, and two TIMs layers. The secondary heat dissipation path goes through the PCB layer below the die to the Motherboard.

ease heat dissipation and facilitate the mounting of the aluminum heat sink. From physics first principles and applying the Fourier heat-equation to the silicon and metallic layer, we obtain the following Partial Differential Equations (PDEs):

$$\rho_{Si}c_{Si}\frac{\partial T_{Si}(x,t)}{\partial t} = k_{Si}\nabla^2 T_{Si}(x,t) + q(x,t), \quad \text{with: } x \in V_{Si}, t \in \mathbb{R}_{\geq 0} \quad (1)$$

$$\rho_{Cu}c_{Cu}\frac{\partial T_{Cu}(x,t)}{\partial t} = k_{Cu}\nabla^2 T_{Cu}(x,t), \quad \text{with: } x \in V_{Cu}, t \in \mathbb{R}_{\geq 0}, \quad (2)$$

where $T_{Si}(\cdot)$ and $T_{Cu}(\cdot)$ are the temperatures of the silicon device and the heat spreader defined in the open volumes V_{Si} , V_{Cu} , $q(\cdot) \geq 0$ is volumetric thermal power generated by internal sources (e.g. cores power outputs), and ρ_{Si} , ρ_{Cu} , c_{Si} , c_{Cu} , k_{Si} , k_{Cu} are the density, specific heat and thermal conductivity of the two materials respectively.

An approximated model of the aluminum heat sink and the PCB is presented in Section 3.1 as an extension to the model of the cores and the heat spreader since (i) the time-constants for the heat sink and the PCB are three to four order of magnitude slower than the faster thermal time constant of the Silicon [53] making it less relevant from the controller point of view, (ii) the heat sink thermal models is highly variable, depending on its shape, characteristics, and the variable airflow of the fans, and (iii) the PCB is can not be modeled and the lower surface of the cores can be considered adiabatic [8, 54]. Thus, modeling them with the same precision as the cores will introduce unnecessary complexity in the thermal model. For the same reasons, the model of the




Fig. 4. Example of finite elements spatial discretization, with corresponding lumped parameters model of a single PE. In the model, the green part (below) concerns the core, with P_k being the consumed power generated by the PE, the yellow part (above) concerns the heat-spreader, with T_E being the main dissipation path to the Heat Sink.

interposer layer, required by the chiplet design, is merged with the PCB model.

To obtain a model which can be tractable for control design purposes (control model) the PDEs in (1)-(2) can be “converted” into Ordinary Differential Equations (ODEs) by employing Euler discretization of the derivatives with respect to the spatial coordinates x . Clearly, this operation introduces some degree of approximation. However, resorting to the structural properties of thermal systems, for which the *Maximum principle* [23] holds, it can be shown [8, 53, 56] that the feasibility of the thermal constraints can be guaranteed by discretizing the original PDEs in space and focusing on the hot spots of the die. Therefore, for simplicity we select n_c finite elements (cells) each associated with a power source, i.e. a PE. Multiple elements could be assigned to each power source with a slight increase in the notation complexity, similar to what is shown in Figure 4. Note that each finite element has two thermal states, one associated with the (local) temperature in the silicon die and the other with the local temperature of the metallic heat spreader. The power source P_i associated to each source is described as follows:

$$P_i(t) = \int_{V_{s,i}} q(x, t) dV = h_i(F_i(t), V_i(t), \omega_i(t), T_{Si,i}(t)), \quad (3)$$

where h_i is the nonlinear mapping function. Different models can be used for h_i without significant modifications to the proposed HPC processor model. The power model (4) employed in this work is the most commonly used one [54].

$$\begin{aligned} P_i &= P_{i,\text{stat}} + P_{i,\text{dyn}} \\ &= k_{s_0} + (I_{cc,i} V_i) \cdot \mathcal{K}(T_{Si,i}, V_i) + C_{eff,i} F_i V_i^2 \end{aligned} \quad (4)$$

where $P_{i,\text{stat}}$, $P_{i,\text{dyn}}$ are the static and dynamic component of the PEs’ power consumption, depending on the static current $I_{cc,i}$, and the effective capacitance $C_{eff,i}$ respectively. C_{eff} is correlated to the type of instructions ($\omega_i(t)$) being executed on the PE. $\mathcal{K}(T_{Si,i}, V_i)$ is a non-linear mapping that encapsulates the dependency of the static leakage power to the temperature and the voltage of the component.

Adding the temperature dependency to the power model (i) raises the complexity of the control algorithm since both the temperature T_i and the power consumption P_i of each component are tightly coupled between each other, (ii) creates scenarios where constant inputs (F_i , V_i) and workload ω_i produce variation in the power consumption P_i , and (iii) models the thermal run-away scenario

[58] where temperature and power increase each other in a positive feedback, leading to the irreversible damage of the component.

In this work an exponential relation based on [9, 48] is employed:

$$\mathcal{K}(T_{Si,i}, V_i) = e^{(k_v V_i(t) + k_T T_i(t) + k_{T_0})} \quad (5)$$

where the k parameters are constant and computed on the critical values of V_{MAX} and T_{MAX} .

Denoting with $\mathcal{N}\{i\}$ the set of all neighbors of the element i , then the differential equation associated with a generic element reads as:

$$\begin{aligned} \dot{T}_{Si,i} &= \frac{P_i}{C_{Si,i}} + \frac{T_{Cu,i} - T_{Si,i}}{C_{Si,i} R_{Si,i}^v} + \sum_{j \in \mathcal{N}\{i\}} \frac{T_{Si,j} - T_{Si,i}}{C_{Si,i} R_{Si,i,j}^h} \\ \dot{T}_{Cu,i} &= \frac{T_{Si,i} - T_{Cu,i}}{C_{Cu,i} R_{Si,i}^v} + \frac{T_{Al} - T_{Cu,i}}{C_{Si,i} R_{Cu,i}^v} + \sum_{j \in \mathcal{N}\{i\}} \frac{T_{Cu,j} - T_{Cu,i}}{C_{Cu,i} R_{Cu,i,j}^h}, \end{aligned} \quad (6)$$

where $C_{Si,i}$, $C_{Cu,i}$ are the thermal capacitances, while $R_{Si,i}^v$, $R_{Cu,i}^v$, $R_{Si,i,j}^h$, $R_{Cu,i,j}^h$, $j \in \mathcal{N}\{i\}$ are respectively the vertical and horizontal thermal resistances, and T_{Al} is the temperature of the aluminum heat sink. These lumped parameters stem from the aforementioned spatial discretization procedure as shown more in detail in [40] and [53].

$$\begin{aligned} C &= c^{Si/Cu} \cdot h \cdot w \cdot l \\ R^v &= k^{Si/Cu} \cdot \frac{h \cdot w}{l} \\ R^h &= k^{Si/Cu} \cdot \frac{l \cdot w}{h} \end{aligned} \quad (7)$$

Collecting all temperatures in a unique vector $T = (T_{Si,1}, T_{Cu,1}, \dots, T_{Si,n_s}, T_{Cu,n_s}, T_{Al})^T$ we can compactly rewrite system (6) as:

$$\begin{cases} \dot{T}(t) = \mathbf{A}T(t) + \mathbf{B}P(t) \\ T_{Si} = \mathbf{C}T(t) \end{cases} \quad (8)$$

where $P = (P_1, \dots, P_{n_c})^T$, and $\mathbf{A}$, $\mathbf{B}$, $\mathbf{C}$ follow directly. Dynamics in (8) combined with the algebraic power model in (4) describe the overall system thermal and power behavior.

3.1 Thermal Model extensions

The modeling of the aluminum heat sink and PCB is realized by adding a state in the system representing each entire layer. These states exchange heat with the connected blocks through heat conduction, with the addition of a spreading resistance in series to account for the different contact areas, as shown in [53] following the formula in [30]. Albeit being an approximation [46], with this approach we aim at introducing slow time-varying temperature and dissipation shifts to study the impact on the controllers.

To complete the thermal dissipation path, two additional states for the motherboard and the air are added to the thermal model, as well as additional resistances to emulate the presence of the two different Thermal Interface Materials (TIMs) between the cores and the heat spreader, and the latter and the heat sink. The vectors and matrices of the model (8) are extended with the inclusion of the described four states.

3.2 Modeling Real Actuators

As shown in Figure 2, the control outputs are the Frequency F and the Voltage V . Most of the control challenges proposed in this work derive from the non-idealities of these two actuators.

The Frequency and Voltage variables are not independent. A Field Effect Transistor (FET)-based digital circuit requires continuously toggling the FETs's state, changing its gate voltage below or above a certain threshold. The rate of speed at which this change can happen (the device *switching speed*) contributes to the maximum frequency at which that digital circuit can run operations. Toggling the FET by changing its gate voltage requires charging or discharging its transistor capacitance. Higher voltages result in a faster slew rate when charging and discharging, allowing an increased switching speed. Additionally, the further away the gate voltage from the threshold voltage, the lower the resistance of the FET's conducting channel, resulting in a lower RC time constant and thus a quicker realization of the logic state. For these physical reasons, maximum frequencies are bounded by the currently supplying voltage. Reaching higher frequencies requires increasing the voltage, and the relation between Frequency and Voltage is sub-linear [9, 43].

Another non-ideality that is common in multi-cores processors design is that Frequency and Voltage actuators may be shared among several PEs. This is often the case for the Voltage, which is shared among several cores that are thus divided into *Voltage Domains*, while it is more common to find PEs with exclusive PLLs or frequency dividers. Requiring multiple cores to share the same voltage introduces a coupling effect in the output of the controller, which may result in degraded performance.

Indicating with $\mathcal{D}_j$ the j^{th} power rail that supplies the component c_i , and omitting the time dependency for ease of reading, we can describe the two above restrictions as:

$$F_i \in [F_{\min}^S, F_{\max}(V_{\mathcal{D}_j}, T_i)] \quad \text{with } c_i \in \mathcal{D}_j \quad (9)$$

where $F_{\min}^S$ is the minimum core frequency of the system.

Frequency and Voltage are typically generated by Phase Locked Loops (PLLs), Frequency Locked Loops (FLLs), or frequency dividers and Voltage Regulator Modules (VRMs), respectively. These components generate discrete values and introduce delays (*transition times*) when changing from one discrete value to the other, requiring either to halt the processor or to keep a minimum value for the transition time, decreasing the application execution performance [42].

In this paper we model these delays with the following assumptions: (i) we assume two PLLs for each frequency output, thus when one is changing to a new value the other can maintain the current frequency without halting the processor, (ii) voltage transitions require waiting at a saturated frequency the completion of the transition, thus we can assume a constant maximum delayed response in the changed output. Although the voltage transition is completed in less than $50\mu\text{s}$ which is faster than the control loop [42], it is essential to model these delays as they add to the whole controller response latency, allowing us to study more in detail the impact of the control step periodicity on the control performance (see Section 4.2 - Controller delays).

3.3 Exponential Leakage Analysis

Figure 5 shows the comparison between the exponential leakage power model (EXP-LP) (5) and the same model without exponential relation (L-LP). In the first plot, we observe the interval of values that characterize the exponential relationship of EXP-LP to V and T , and the gap with the flatter magenta plane (L-LP model), with a power consumption that can be up to 10 times higher at high voltages and temperatures. The top view of the graph in the bottom left shows that for most operating points the EXP-LP model power consumption is less than or equal to the L-LP model, thus the exponential behavior of the static leakage power shows at operating points that




Fig. 5. Multi-view comparison between the proposed model of the leakage power characterized by an exponential relationship with Voltage and Temperature (blue surface), the same model of leakage power with no exponential relation (magenta surface), and the minimum and maximum dynamic power (yellow and red surface respectively).

are compatible with high-performance scenarios.

In the other plots, we can observe that at those operating points, the leakage power becomes a relevant or even predominant fraction of the total power consumption.

3.4 Model Implementation in Simulation

There are two different timing requirements to run the presented model in a simulation: the timing restriction related to the power model (4) and the one related to the thermal model (8). According to [12] the thermal timing constants inside a processor are $\sim 10s$, $\sim 0.1s$, $\sim 1ms$. The power consumption of each component $P_i(t)$ varies according to the workload $\omega_i(t)$ and the frequency of operation $F_i(t)$. Assuming the frequency to be in the GHz scale, and that the workload could be constantly changing, the power consumption $P_i(t)$ would be changing with timings below the ns scale.

Running a simulation with ns steps would be superfluous for two reasons: power spikes between $10\mu s$ and $100ns$ are mitigated by hardware elements of the Power Delivery Network (PDN) [20], and it is extremely difficult to obtain instructions sequences with a granularity of ns since processors inspection tools provided by the manufacturer (as event and performance counters) have generally a sampling time of $1ms$ [27]. Choosing a simulation step t_s that is slower than the frequency F , requires to use an average power consumption $\bar{P}_i(t_s)$ instead of the instantaneous one.

Introducing $\mathcal{S}_i = (C_0, C_1, \dots)$ as the sequence of *instructions-characterized cycles* to be executed on the i -th core², each element of the sequence corresponds to 1 cycle, thus describing not only the instructions sequence but also characterizing instructions IPC, memory stalls, cache misses,

¹1ms is the default tick period of the operating system scheduler, which makes performance monitoring less intrusive if done synchronously with it [9, 52].

²Consider this as a simulation abstraction since the instructions sequence to be executed and the IPC are not generally known a priori.

SIMD/Vector instructions, etc. During each simulation step t_s , assuming the core frequency $F_i(t)$ and the voltage $V_i(t)$ to be constant, a variable number of instructions $\theta_{v_i}(t_s)$ from S_i will be executed, with $\theta_{v_i}(t_s) = F_i(t_s) \cdot t_s$. Assuming the simulation step t_s to be appropriate orders of magnitude faster than the fastest thermal time constant, in each t_s only the $C_{eff,i}$ component of the model (4) will vary depending on the type of workload. Thus computing an effective capacitance density across each t_s , would allow us to compute $\bar{P}_i(t_s)$.

$$\begin{aligned} \bar{P}_i(t_s) &= k_{s_0} + (I_{cc,i} V_i) \cdot \mathcal{K}(T_{Si,i}, V_i) + \overline{C_{eff,i}} F_i V_i^2 \\ \text{with: } \overline{C_{eff,i}} &= \frac{\sum_{i=1}^{\theta_{v_i}(t_s)} C_{eff,i}(t)}{\theta_{v_i}(t_s)} \end{aligned} \quad (10)$$

We chose a simulation step of $50\mu s$, which is two orders of magnitude faster than the fastest timing constants and in the range where the Power Delivery Network mitigates power spikes. We introduce parameter variation in the coefficients of the power model (4) and thermal model (8) to encapsulate the manufacturing difference and process variation from core to core [5, 9].

Finally, we model white noise at the output of the thermal model T_{Si} to emulate the cores Process, Voltage, Temperature (PVT) sensors [22].

4 CONTROL PROBLEM

4.1 Formulation

From a thermal perspective, each component of a CPU has a critical thermal limit T_{CRIT} causing irreversible damage if exceeded. High temperatures lead to faster aging as well as performance and reliability degradation [16], modeled through a thermal limit T_L . Compared to T_{CRIT} , T_L is a softer restriction that, if exceeded temporarily and infrequently, does not cause damage.

Another type of physical limitation is given by the maximum power that can be supplied to the processor. As the effects of Dennard scaling came to an end in the first decade of 2000, the power density constantly increased at each technology scaling, limiting the power and the current in today's processors. Pads in the package are shared among power delivery networks, memories, and I/O interfaces, generating a trade-off between peak current and data bandwidth. At the board level, VRMs capacitors and inductors become dominant. As a result, modern processors' clock frequency is bounded by the power that can be provided to the processor itself in a given instant and not always by its absolute and attainable physical speed. For this reason, system-level and voltage-domain-level power consumption limits need to be enforced.

As described in Section 3, the temperature T_i and the power consumption P_i of a component are coupled in a positive feedback, as they increase when each other increases. High P_i and/or T_i cannot be sustained for a long period of time, as they can lead to an uncontrollable system condition referred to as *thermal runaway* [58]. This detrimental behavior is caused by the exponential relationship between the leakage power and the temperature, as described in Section 3.3.

In addition to the thermal and power limitations due to the physical nature of the components, there are additional time-variable power and thermal limits in the form of requests from external agents, such as off-chip board and node management controllers, or the OS. Assuming the requested limits are lower than the physical ones, we consider only the former in the control problem formulation.

The performance objectives of the control algorithm are (i) the time of completion t_{exec} of the application running on the HPC chip, and (ii) the energy efficiency as in the energy E_{tot} required to complete the execution [9, 19]. E_{tot} can be computed as the integral of $P_T(t)$ over t_{exec} . $P_T(t)$ is the sum of all elements of the chip consuming power. Note that from model (4), $P_T(t)$ has a static

and a dynamic contribution. Hence, the execution speed plays a role in the energy efficiency. t_{exec} depends on the PEs running frequency, and its Instructions per Cycle (IPC).

Both the performance objectives described above depend on an exhaustive knowledge of the running application. For security and privacy concerns, the integrated LLC has no privileged access over the latter. Additionally, energy savings depend on other external components (GPU, cooling system, power supply, ...) over which the LLC has no control, and, if the application execution is spread on different systems, the LLC has no global access on those systems too. The choice of the optimal operating point and the desired trade-off between t_{exec} and E_{tot} is therefore left to the HLCs. The LLC's objective is to apply the received operating point while enforcing the physical and requested thermal and power constraint described above within the required timing restrictions. Being implemented as a dedicated unit tightly coupled to the system under control, the LLC can perform low-level sensing and actuation at a faster pace than HLCs. Reductions in the operating points performed by the LLC due to thermal and power capping should still reflect the performance objectives dictated by the HLCs.

Provided the description above, the control problem is formulated as follows:

$$\begin{aligned} & \min_{F_a} \left| F_T(t_k) - F_a(t_k) \right|_{\mathcal{R}}^2 \\ & \text{subject to : } \sum_{i=1}^{n_c} P_i(t_k) \leq P_B(t_k) \\ & \sum_{i=1}^{n_j} P_i(t_k) \leq P_{\mathcal{D}_j}(t_k), \quad i \in \mathcal{D}_j, \quad j = 1, \dots, n_d \\ & T_i(t_k) \leq T_L(t_k) < T_{\text{CRIT}}, \quad i = 1, \dots, n_c \end{aligned} \tag{11}$$

where $F_a = (F_{a1}, \dots, F_{anc})^T$ are the applied Frequencies to the n_c cores, $F_T = (F_{T1}, \dots, F_{Tnc})^T$ are the given Target Frequencies as operating point, and $\mathcal{R} \in \mathbb{R}^{n_c \times n_c}$ is a symmetric positive definite weight matrix and $|\cdot|_{\mathcal{R}}$ is the corresponding norm. P_B is the requested total power budget, n_d is the number of the $\mathcal{D}_j$ voltage domains such that $\sum_{j=1}^{n_d} n_j = n_c$. Finally, T_i is the i^{th} core temperature.

The computation of the total power consumption as the sum of the power consumption over all the PEs (11.2) is not fully accurate as it neglects other entities such as DRAMs, on-chip high-bandwidth memories (HBMs), I/Os controllers, and the case of multiple chiplets. In this work, we focus on the PEs control in a single chiplet scenario only.

4.2 Challenges

There are numerous state-of-the-art solutions to address the control problem (11) [19, 49, 60]; however, many of these solutions do not fully account for all the control challenges involved.

First, the MIMO nature of the problem and the non-linearity of the controlled system, as shown in equations (4) and (5). Then, the tightly coupled control objectives, the coupling, constraints, and discretization of the control outputs, and the high-amplitude and high-frequency unpredictable noise originated from the unknown workload $\omega(t)$. Finally, the algorithm has to be simple enough to execute with a frequency in the order of $2 - 10\text{kHz}$ on an embedded microcontroller. This frequency accounts for both the requirements given by the limitation of the PDN and the faster thermal time constant (as described in Section 3.4 and in Section 4.2 - Controller delays), allowing a consistent control of the temperature and power capping according to the most used control rules [24].

Controller outputs. Figure 6 shows the progression of the power output of a PE, according to the model of a real actuator. In the left figure, we observe that, except for the initial values, the




Fig. 6. The PE power consumption $P(t)$ according to the model (4) in relation to frequency F (left) and the voltage V (right) at 75°C . Each graph only shows operating points that are feasible w.r.t. (9). Different lines represent the non-plotted variable, i.e. the voltage in the left graph, and the frequency in the right one.

power gap between different voltage levels is larger than the line slope along frequency values. The power gap along the maximum $F - V$ line (blue) increases by more than $3\times$ at 2.2GHz and $5\times$ at 2.7GHz . This means that, once a voltage level is chosen, changing frequency within that level (up to the maximum allowed frequency relative to that voltage level Section 3.2) generates little power (and thus thermal) differences. Instead, changing the voltage level causes larger power and thermal deviations. For example, along the max $F - V$ blue line in Figure 6, increasing the frequency of 50MHz from 2.85GHz to 2.90GHz at 0.9V will produce a 62.7mW power increase, while the same frequency increment, from 2.90GHz to 2.95GHz (involving an increase in the voltage level since 2.9GHz is the maximum allowed frequency for 0.9V), produces a power increase of 740.7mW , which is more than $10\times$. The control algorithm should consider this difference when choosing the operating point of each PE.

The example above is made more complex by the shared voltage among cores in the same domain, which introduces an additional coupling in the outputs of the controller (Figure 2). Increasing the voltage consists in increasing the power consumption (and temperature) of all the cores of the domain. Decreasing the voltage means capping the maximum frequency, and thus the performance attainable by the cores in the domain that were not hitting any power or thermal limit.

The control algorithm should consider this coupling when distributing the power or selecting the frequency of each core. This is well illustrated in the right-hand side of Figure 6. For example, let's consider the case where core A is capped at 3.8W for thermal reasons, running at the operating point $\{2.75\text{GHz}, 0.85\text{V}\}$. If, for external reasons (such as another core requesting a higher frequency), the domain's voltage is increased to 0.95V , the core A's frequency would have to be reduced to 1.85GHz to retain the same power output. This will result in a performance loss due to the frequency being decreased.

Consider now the opposite case in which core A has thermal headroom and could run at its maximum speed. If a different core B in the same voltage domain has to run at a lower voltage to

dissipate some previously acquired heat³, core A would be limited to the same maximum frequency as core B, and the power budget would not be used effectively (under-utilization).

Due to the coupling, a controller actuating a voltage increase should take into account whether the corresponding ΔP may incur power budget violations or high-temperature overshoots, requiring the thermal control to be activated to reduce the frequency (and thus performance).

Disturbances. There are three different types of disturbance in a typical CPU control scenario: instruction variability, process variation and drift, and sensor noise. The common control strategies available in the literature are able to tackle the magnitude of sensor noise, but such disturbance prevents precise power/workload identification at higher frequency domains.

Parameters deviation in modern and old technology nodes are related to the manufacturing flow, especially the front-end-of-line (FEOL). Previous works report that they can cause variations of up to 10% in the current flowing through the manufactured devices [5, 11, 59]. Parameters can also slowly change during utilization due to aging. These deviations impact the power consumption and thermal response of the PE. Control algorithms should take into account this differentiation and not treat each PE equally.

Instructions variability, even if smoothed by the power delivery network, is a constantly varying noise input to the controlled system. This disturbance is unknown, varying with a frequency faster than the control algorithm's one, and accounts for a significant part of the power consumption $P(t)$ as shown in Figure 5. The control algorithm can see power consumption fluctuation of several watts at each iteration with no changes to the input values of $F(t_k)$ and $V(t_k)$. Therefore the control action should provide a good balance between reactivity and filtering.

Exponential relation of leakage power and temperature. The exponential relation between the leakage power and the temperature links the two objectives of the control problem (the power and the temperature) to each other. This relation amplifies the non-linearity already lying in the power map (4), creates positive feedback, and is the main cause of the thermal runaway scenario described in Section 4.1. Furthermore, this coupling hinders achieving a steady state, considering the disturbances present in the system that can continuously produce power spikes and oscillations. It also further complicates the challenges and coupling effects described in Section 4.2 - Controller outputs.

Power and thermal oscillations. Non-linearities, output discretization, and output coupling produce high-amplitude power consumption steps, that can cause significant oscillations around the thermal and power capping set points. For this reason, these set-points often cannot exactly be met. These effects are exacerbated when dealing with an increasing amount of cores in the same shared voltage domain. Furthermore, in the latter scenario, the thermal coupling of neighboring cores worsen the oscillations.

High-amplitude oscillations violate thermal and power capping constraints, decrease performances over time, and stress the hardware, causing it to age faster.

Controller delays. To be consistent in the output generation, digital discrete controllers such as the LLC apply the computed output at the beginning of the next periodic execution [24], thus adding a delay of $1 t_s$, with t_s being the controller execution periodicity interval.

In the described control problem, there is an additional source of incremental delay given by the workload ω . The workload directly affects the instantaneous power consumption of the PEs

³From the left-hand side of Figure 6, a PE running at the minimum frequency (0.4GHz) and at the maximum voltage, generates a considerable power consumption (5.52W). This is similar to the power generated by the same PE at {3.0GHz, 0.95V}. This means that, even when running at minimum frequency, a core may need to reduce its voltage level to reduce its temperature.




Fig. 7. Five iterations of the test presented in Bambini *et al.* [7] with the additional thermal and power model features introduced in Section 3.1. From the left-hand side, (1) the original test, (2) the addition of the exponential leakage relation (3.3), (3) the addition of Dynamic Voltage Scaling (DVS), (4) the addition of 3 shared voltage domains, and (5) all the previous together.

(4), but it will only be measured at the following controller interval. If the workload has a sudden change in its average value, which is a common event, the controller response will have a delay of $(1 + \nu)t_s$, where $\nu \in (0, 1)$ is the latency introduced by the workload variation.

This delay is an additional cause of power and thermal oscillations. To reduce the thermal oscillations, a controller interval t_s less than at least half the faster thermal time constant is required.

To reduce power oscillations caused by workload variation instead requires PE's hardware power management features added by the manufacturer. In the past, Intel was dealing with this problem by forcibly restricting the frequency when certain high-power workloads were detected [52], new designs include microarchitecture mechanisms to delay high-power instruction dispatching, spreading the power increment into multiple t_s [2].

5 CONTROL ALGORITHMS

In Section 5.1, an extended and improved version of the control algorithm proposed by Bambini *et al.* in [6] and evaluated in [7] is presented to account for the not-idealities in the control problem introduced in this work. The original version is referred to as *Baseline Algorithm (BA)* and the enhanced version as *Enhanced Baseline Algorithm (EBA)*. A comparison between the performance of the BA and EBA against the *industrial SotA* is then carried out. As described in the related work section, this is based on a Voting Box Algorithm (VBA). Results will show that the EBA underperforms the industrial SotA when the not-idealities are present. In Section 5.2 a Fuzzy-inspired Iterative Control Algorithm (FCA), the proposed algorithm designed to overcome the limitations of the EBA, is described.

5.1 Addressing the non-idealities in the original algorithm

The BA has the following key elements: (i) an initial conversion step (*Conv2P*) in which F_T and ω are used to compute an estimation of each core's power consumption P_{est} , (ii) a cascade of two blocks to perform the power capping and the thermal capping respectively, that outputs the reduced power P_a , and (iii) a final conversion step (*Conv2F*) to compute from P_a the frequency setting for the core F_a .

Figure 7 shows the comparisons of BA and EBA against the VBA in a simulation of a computing tile of 9 PEs simulating a workload lasting 2s. Each of the subplots in Figure 7 reports on the y-axis the relative speedup of each PE (reported on the x-axis). Different colors are associated with different types of workloads: high-power workloads are red (core 1, 3), low-power are blue (core 2,

4), and variable workloads are green (core 5-9). PEs target frequency is set to the maximum to stress the control algorithm in the worst capping scenario. Each subplot introduces in the simulation environment one of the not-idealities described in Section 3. At the same time, we add a different control algorithm feature implemented in the EBA.

Figure 7.1 shows that the more accurate thermal model introduced in this work with the addition of the secondary thermal path and other thermal dissipation elements leads to similar results to the ones obtained in [7].

In Figure 7.2, the exponential relation between leakage power and temperature is introduced. With a fixed high voltage, the power consumption $P_i(t)$ and the temperature $T_i(t)$ grow faster, affecting the control quality. In Figure 7.3 DVS is enabled, and the EBA is switched in place of the BA in the comparison with the VBA. As presented in Section 4.2, different voltage levels have severe variations in the power consumption (and temperatures), more than frequencies. In BA, the voltage has to be chosen in advance since the first step of the algorithm involves the conversion of the input variables into $P_{est,i}$ (Conv2P). In the scenario of a voltage domain per each core, this issue could be solved in Conv2F by selecting a pair of F_i and V_i according to a look-up table. Nevertheless, due to the non-idealities of the actuator described in Section 3.2, such implementation incurs an over-conservative policy in the presence of shared voltage domains. Different cores may have their power reduced for different reasons, and the choice of a voltage *a-posteriori* could severely affect the performance objectives. As shown in section (4.2), increasing the voltage of the domain can cause a performance drop to all the other cores of the domain. A better power distribution would balance the performance of all the cores according to the final voltage choice.

To address such over-conservative behavior and prevent performance losses due to ineffective DVS, a moving average filter for the voltage selection is introduced in the EBA. After each iteration t_k of the control algorithm, the difference between the target frequency $F_T(t_k)$ and the applied frequency $F_a(t_k)$ is added to an accumulation value $F_{MA}(t_k)$ through a forgetting factor λ_{MA} . With several iterations, $F_{V-MA} = F_T - F_{MA}$ should be tracking the average frequency chosen by the control algorithm over a period $\mathcal{T}_{MA}$ dependent on the choice of λ_{MA} . In the Conv2P part, $F_{MA}(t_k)$ is used to select the initial shared voltage of the domain to have a more realistic evaluation of the power output of each component. However, the target frequency $F_T(t_k)$ is still used in the computation of the estimated power P_{est} to allow the control algorithm to increase the applied frequency when thermal and power headrooms are available. At each iteration, the EBA executes:

$$F_{V-MA}(t_k) = F_T(t_k) - F_{MA}(t_{k-1}) \quad (12)$$

$$V(t_k) = f_V(F_{V-MA}(t_k)) \quad (13)$$

$$P_{est} = h_{est}(F_T(t_k), V(t_k), \dots) \quad (14)$$

$$F_{MA}(t_k) = (1 - \lambda_{MA})F_{MA}(t_{k-1}) + \lambda_{MA}(F_T(t_k) - F_a(t_k)) \quad (15)$$

where $f_V(F_i)$ is a step-wise monotonically increasing function that, given a frequency F_i , computes the minimum voltage such that (9) is fulfilled.

Equation (15) can be thought of as a discretized version of a low-pass filter where the forgetting factor λ_{MA} is equal to the ratio between the controller sampling time T_s and filter time constant τ . The value of τ should be chosen to reject small system state perturbation caused by workload variation and intercept the main tendencies. In this work the value of τ is chosen as $\tau = 1.25 \cdot 10^{-2}$. This method decreases the performance during transients after a change in one of the given set points or system conditions, depending on the choice of the forgetting factor. A strategy to detect application phases could be employed to partially avoid this performance reduction, by resetting $F_{MA}(t_k)$ after a different phase is detected or a set point change is required from HLCs, significantly decreasing the “adaptation time”. Observing Figure 7.4, the addition of shared voltage domains and

a variable voltage lead to results similar to the original test from [7] shown in Figure 7.1.

A second necessary modification to the BA is due to the larger difference between the measured and computed power P_{est} , originating from the exponential relation highlighted in (5). This can lead the BA to miss the power budget set point by a large and variable margin (positive or negative). This approach is in contrast to classical predictive approaches based on an a priori knowledge of the power model [8], but this would require a non-negligible effort to study all different operative points. Our solution is simpler and employs a moving average adaptation filter. The difference between the sum of P_{est} and the measured power is added in the power distribution block of the algorithm to take into account the power model difference through a moving average filter to avoid high-frequency noise and oscillations.

A third observation is how to choose the voltage that is shared among all the cores in the domain. In EBA, each core requests a voltage level based on the minimum voltage required to run at the given frequency (in this case, the $F_{V-\text{MA}}$). SotA in industry chooses the maximum voltage, i.e., the voltage level that is compatible with the highest frequency requested in the voltage domain, according to (9). This approach causes a larger performance drop if frequencies in the same domain have larger value differences. This behavior is shown in Figure 7.5 where the EBA's results are worse than SotA.

Other control decisions to be investigated are (i) the selection of a voltage as a percentile of the casted vote (e.g. 90%), instead of the maximum, or (ii) running an optimization step that computes the voltage that will cause the minimum sum of performance differences according to an 'importance matrix' $\mathcal{R}$.

We can conclude that the EBA approach is too simplistic to cope with the power and thermal management of many-cores systems when non-idealities are considered. In the next section, we propose a different algorithm to tackle them.

5.2 Proposed Fuzzy-inspired Iterative algorithm

As from Figure 7.5, there are some scenarios where the EBA does not perform as well as the reference. To further improve control quality, the control algorithm was redesigned retaining its advantages while addressing at the same time some of the criticality observed in the previous section.

The algorithm is still based on a cascade design with power as the main control variable. The problem with the *a priori* choice of the shared voltage in the domain is solved by employing an iterative solver method in the Conv2F part, converting the power into the pair of frequency and voltage.

The SotA approach of selecting the maximum voltage can be expressed as $V_{\mathcal{D}_j} = \max(f_V(F_{i \in \mathcal{D}_j}))$. Since $f_V(F_i)$ is step-wise monotonically increasing, the right term can be rewritten as:

$$V_{\mathcal{D}_j} = f_V(\max(F_{i \in \mathcal{D}_j})) \quad (16)$$

Denoting h_{est} as mapping (3) of the control algorithm to compute P_{est} and P_a as the output of the control algorithm, a system of equations for each domain $\mathcal{D}_j$ can be derived where P_a equals h_{est} combined with (16). This way, the frequency-voltage relation (9) is directly enforced in the conversion step.

$$\begin{cases} P_{a,1} = h_{\text{est}}(F_1, f_V(\max(F_{i \in \mathcal{D}_j})), \dots) \\ \vdots \\ P_{a,n_j} = h_{\text{est}}(F_{n_j}, f_V(\max(F_{i \in \mathcal{D}_j})), \dots) \end{cases} \quad (17)$$

Assuming $\forall i = 1, \dots, n_j$

$$\begin{aligned} P_{a,i} &> \min[h_{\text{est}}(F_{\min}^S, V_{\mathcal{D}_j}, \dots)] \\ P_{a,i} &< \max[h_{\text{est}}(F_{\max}^S, V_{\mathcal{D}_j}, \dots)] \end{aligned} \quad (18)$$

a solution to the system of equations (17) is guaranteed to exist.

To solve these j non-linear systems, the global Newton-Raphson iterative method could be employed, using the Softmax function $\mathcal{S}_\alpha$ [25] as an approximated $\max()$ function that is differentiable. Even with hardware-optimized algorithms, the time to solve the global N-R with a reasonable tolerance is not feasible for a real-time control solution running on an embedded LLC with tight periodicities such as [39].

Considering that the solutions' accuracy is bounded by the precision of the system actuators (generally not smaller than $5 \cdot 10^{-3}$), the bisection method [15] was chosen instead. The bisection takes more iterations to converge, but each iteration is faster compared to Newton-Raphson. Furthermore, it does not require inverting the Jacobian and its execution time is less dependent on domain dimension. From a deployment perspective on an embedded LLC, this means reducing the execution time, hence being able to execute within the control period.

Since the bisection does not need derivatives, the $\max()$ function can be employed. Furthermore, the function f_V can be upper-bounded by a non-linear polynomial function, such as the second order function $f_V(F_i) = F_i^2 + k_1 F_i + k_0$. The coefficients of this function can be found *a priori* since they depend on the F-V relation (9). This solution still didn't address the problem of the output discretization as both the F and V values are considered continuous.

However, using an iterative method for the Conv2F part instead of the F_{MA} in conjunction with the original PID control algorithm, yielded poor results. The PID is a linear time-invariant controller. Albeit the thermal evolution of the system can be described by a linear relation, as non-idealities and coupling such as the ones described in Section 4.2 are introduced in the model the limitations of the PID emerge. In particular, the discrete coupled outputs which are exponentially related to the temperature state, require a PID tuning pushed towards a fast response to avoid the thermal runaway scenario. These types of PID tuning also generate high-frequency output oscillations which cause control performance losses, and these oscillations were more frequently present in combination with the iterative method. For this reason, the thermal capping control was changed to a heuristic LUT-like control inspired by the Fuzzy control theory.

The fuzzy control theory is a simple and effective way of turning symbolic decision tables into control laws. In practice, fuzzy logic consists of three main steps: (1) the control inputs are converted into fuzzy quantities (i.e., sets) through a process called *fuzzification*; (2) logic rules are applied to decide the control action, and (3) the latter is finally converted into a quantitative control output value in a process called *defuzzification* [38]. Fuzzy control can be applied to non-linear systems, it does not need a model, it is simple to understand and compute, and it is modular and easily modifiable. Hereafter, we refer to the fuzzy-inspired control with the iterative operating point selection as FCA.

Due to the capability of fuzzy logic to directly tackle non-linearities, the thermal control can be directly applied to the frequency F and voltage V values. For this reason, in FCA the thermal capping block is moved before the power distribution block. With this structural modification, already distributed power from the power capping block is not wasted to a later potential further power reduction for thermal capping reasons. The control structure in BA was motivated by the power affecting the temperature, and our proposed control mimicked that structure. But in the updated model, the temperature is affecting the power too, making the previous statement no longer valid.

In the FCA, a *fuzzy state* representing the reduction in frequency and voltage due to thermal

$\Delta T/T[^\circ\text{C}]$	< 45	$45 - 65$	$65 - 80$	$80 - T_L$	$\geq T_L$
< 0	2	2	1	0	-1
$0 - 0.5$	2	1	1	0	-1
$0.5 - 1.0$	1	1	0	-1	-2
$1.0 - 2.0$	1	0	-1	-2	-3
≥ 2.0	0	0	-2	-3	-4

Table 2. LUT table to determine the increment/decrement of the fuzzy state on each iteration based on the derivative of the temperature $\Delta T(t_k)$ and the temperature value $T(t_k)$.

capping reasons is assigned to each PE. A look-up table (LUT) is created by dividing PEs current temperature $T_i(t_k)$ and its derivative $\Delta T_i(t_k)$ into regions relevant to the thermal control, and is populated with values to either increase or decrease the fuzzy state. The derivative is approximated as the difference between the current temperature and the temperature of the previous s -th step $\Delta T(t_k) = T(t_k) - T(t_k - s)$, as a means to filter the sensors' noise. The choice of s depends on the ratio between the fastest thermal time constant of the system and the execution period of the thermal part of the control algorithm.

Each negative integer value of the fuzzy state decreases half of the frequency step between two consecutive voltage levels, meaning that each even negative fuzzy state value reduces one voltage level. This configuration could also incorporate a frequency turbo boost feature through positive fuzzy states. Since we do not analyze the turbo capabilities in this work, the fuzzy states are saturated to 0 to maintain a balanced comparison with the other algorithms.

The choice of the number of columns and rows of the LUT depends primarily on the memory and computation constraints of the LLC, and secondary on the desired quality of the thermal control. The choice of the thresholds values of the derivate mainly depends on the thermal time constants of the system, while the choice of the temperature thresholds mainly depends on the slope of the P-T leakage relation (Figure 5). The LUT values can be chosen by discretizing a surface describing a trade-off between performance reduction and control performance. The LUT used in this work is showed in Table 2.

In FCA the power distribution part is also slightly modified to reflect the presence of voltage-sharing domain configurations. The BA power distribution algorithm is applied to each domain rather than each core. A linear combination between the maximum temperature and the average temperature of the cores in each domain along with the sum of all domain's cores P_{est} are used as inputs of the heuristic power capping algorithm. Then, for each domain, the power is distributed based on the workload ω_i of each core, trying to favor high-demanding cores.

6 EVALUATION METHODOLOGY

A fair comparison among control algorithms is challenging to achieve, for multiple reasons. First, the control quality is tightly coupled with the microarchitecture of the system to be controlled (i.e. the HPC processor). Second, there is no standardized metric for this comparison, since software benchmarks compare the overall HPC processor performance and are not focused on the controller alone. Finally, it is difficult to find updated information, algorithms, hardware descriptions, etc. that are covered behind industry secrets.

The most popular industry-standard control is the Voting Box Algorithm (VBA) [14, 27, 47], where several operating points are computed independently one from the other, and then the minimum is selected as the output of the controller. One of the available resources for comparison is the IBM OCC control algorithm from the IBM Power9 chip, whose software code is fully open-source [29]. We use this baseline for a reproducible comparison.

Including in the comparison the HLCs presented in the Related Work would lead to misleading

results. As described in Section 2, HLCs and LLCs are designed to work together and they are developed for their specific objectives, execution scope, and time domain. Despite being able to leverage higher computational capabilities, HLCs would be outperformed by LLCs in the comparison due to their higher controller step time, their assumptions and abstractions of the system (in particular, the assumption that an LLC beneath will be controlling higher dynamics), and the absence of significant features, such as receiving an operating point or capping the power.

For these reasons, the test will include a comparison between the reference VBA with both the EBA discussed in Section 5.1, and the FCA introduced in Section 5.2, to observe the benefits and the drawbacks of each design.

We identify three different workload scenarios to investigate the response of each control algorithm: (i) the MAX-WL test, where all PEs execute a vector workload consuming the maximum power and the target frequency is set to the maximum frequency, (ii) the MULTI-WL test, where some PEs execute vector workload at the maximum frequency (3.45GHz), others perform a mix of floating point and int workload at a lower frequency (2.7GHz), while the rest of the cores (the majority) idle at the minimum frequency (0.4GHz); each voltage domain includes at least one PE from both the first two groups to stress the voltage coupling control constraint, and (iii) the CLOUD-WL test, where random and varying workload is executed on all cores, while the maximum frequency is always required. To see the control algorithm fully in action, the power budget of the processor is changed four times during each test, with both increasing and decreasing patterns.

These tests are evaluated on three different variants of the thermal model (8) to stress the adaptability of each control algorithm to different cooling conditions. We identify a water-cooled model (WATER) where the temperatures of the cores are more homogeneous, an air-cooled model (AIR) where the temperature of the cores has a Gaussian-shaped spatial distribution, and a rack air-cooled model (RACK) where air is pushed from the front and exhausted on the back, generating a column-wise spatial distribution of the temperatures modeled with an additional column-wise temperature coupling among cores.

These tests are also evaluated on four different voltage domain configurations to analyze the impact of voltage coupling and the adaptability of each control algorithm. The configurations are: (i) all cores share a single voltage domain (1-D), (ii) the cores are equally divided into 4 domains (4-D), (iii) the cores are equally divided into 9 domains (9-D), and (iv) each core has its own voltage (A-D), to simulate the case without voltage coupling among cores.

The main metrics to be evaluated when comparing tests' results can be directly derived from the control objective (11). These can be translated to:

- Thermal capping:
 - Exceeded value: maximum exceeded temperature [$^{\circ}\text{C}$] compared to the capping target
 - Total time: percentage of time [%] exceeding the capping target.

Since the thermal threshold is taken below the critical thermal limit (4.1), what would cause the system to fail is the maximum exceeded temperature. For aging and degradation, the metric to consider is the amount of time that the threshold is exceeded.
- Power capping:
 - Exceeded average value: the ratio between the average exceeded power with its power budget target [%]
 - Total time: percentage of time [%] exceeding the capping target⁴.

It is important to know the average exceeded power and time since spikes are tolerated and smoothed by hardware features [20].
- Target Compliance:

⁴We consider only the time when such excess is possible. For example, when there is no power budget, it cannot be exceeded.




Fig. 8. The thermal and power evolution of the proposed comparison test. In particular, this is an AIR model with a 4-D configuration, executing the CLOUD-WL and it is controlled by the FCA algorithm. The first three rows show the temperature evolution of the thermal dissipation paths (1), the heat-spreader hotspot blocks (2), and the cores (3). The last row shows the total consumed power measured. The red dashed lines represent the thermal and power limits.

- The L2-norm⁵ of the difference $\Delta F_{pe} = F_T - F_a$, normalized by the number of cores and control steps [GHz/s];

Other figures for the comparison are the **average** and **min execution progress** (AV/MIN- Wl_p) across all cores, as the amount of workload executed [%] with respect to the case where there is no control action. This serves as a direct index for the impact of the control algorithm on the application execution, with the MIN- Wl_p metric describing the most penalized core.

The power and thermal model have a randomly generated (but constant among all tests) set of parameter uncertainties, and all temperature sensors have a white noise with an amplitude of 1°C as it is commonly found on PVT sensor's specifications [22]. The tests run for 2s using a MIL framework consisting of an emulated computing chiplet with 36 cores. All the tests are executed for 10 iterations with different initial conditions, to render the results independent from the initial state of the system. An example of the test is shown in Figure 8 as a reference for the power and thermal evolution.

7 EXPERIMENTAL RESULTS

Results data are aggregated into figures according to workload type (10), Domain number (11), Model type (12), and tests iteration (13). An additional figure with violin plots (9) serves to describe the result distribution among all tests.

7.1 Simulation Results

Thermal Capping Metrics. FCA outperforms both EBA and VBA in both thermal capping metrics, and has the least variation across tests, suggesting that the proposed fuzzy-inspired

⁵Assuming $\mathcal{R}$ in (11) as an identity matrix, meaning all the cores have the same importance.




Fig. 9. Overall distribution of the control performance metrics plotted per algorithm. In the first row are the Thermal metrics, in the middle row are the Power metrics, and in the last row are the Target Compliance and execution progress metrics.

thermal design is an improvement over the other PID controllers. It achieves a maximum exceeded temperature that is 2.44 and 5.52 times lower than EBA and VBA respectively, while having more than 90% lower exceeding time.

EBA has the highest average exceeded time (18.82%), but with a low maximum exceeded temperature across tests (1.94°C on average). This means that a lower thermal margin of about 2°C could be defined to remove almost all the exceeded time. Furthermore, EBA performs better as the number of domains increases. This is probably a heritage of the algorithm being developed with no voltage coupling constraints.

VBA has the worst thermal capping performance among the three analyzed algorithms. It achieves the highest maximum exceeded temperatures (which are dangerous for the hardware) and struggles particularly in the RACK model and CLOUD-WL tests where it reaches an exceeded maximum temperature of almost 20°C with two spurious results over 30°C. The larger the voltage domains are, the worse it performs, indicating difficulties in managing voltage couplings.

Power Capping Metrics. In the power capping metrics, the FCA has a high average exceeding power (13.98%), but with a low exceeding time (3.25% on average). This suggests that the power capping performance degradations are mostly happening on target or application phase changes. The average exciting power distribution has a larger deviation than the other two algorithms and the exceeding time has a large number of outliers which originated from the MULTI-WL tests as Figure 10 suggests. Unusually, the higher the number of domains, the worse the average exceeding power performance gets, suggesting that the problem may be rooted in oscillations.

VBA has a similar behavior as FCA, with the worst average exceeding power (15.14%) and the lowest exceeding time (1.76% on average). The average exceeding power has a large number of outliers, reaching up to 105.84% of the power budget, while the exceeding time is the most consistent among the three algorithms. These outliers are originated from the MULTI-WL tests.

EBA has an opposite behavior to the other two algorithms, with the best and most consistent




Fig. 10. Control performance metrics plotted per workload type (vector -MAX-WL-, combination of vector, INT/Float, and Idle -MULTI-WL-, and random varying workload -CLOUD-WL-) with their relative error bars. In the first row are the Thermal metrics, in the middle row are the Power metrics, and in the last row are the Target Compliance and execution progress metrics.




Fig. 11. Control performance metrics plotted per number of domains (one -1D-, four -4D-, nine -9D-, and one domain per core -AD-) with their relative error bars. In the first row are the Thermal metrics, in the middle row are the Power metrics, and in the last row are the Target Compliance and execution progress metrics.

average exceeding power (3.74%), and the worst exceeding time (32.84% on average) with the larger deviation.




Fig. 12. Control performance metrics plotted per Thermal Model type (water cooling -WATER-, air cooling -AIR-, and horizontal rack cooling, -RACK-) with their relative error bars. In the first row are the Thermal metrics, in the middle row are the Power metrics, and in the last row are the Target Compliance and execution progress metrics.




Fig. 13. Control performance metrics plotted per tests iteration with their relative error bars. In the x-axis are indicated the average initial temperatures of the system. In the first row are the Thermal metrics, in the middle row are the Power metrics, and in the last row are the Target Compliance and execution progress metrics.

Target Compliance and Execution Metrics. On average, the FCA has the best target compliance across all tests, with the exception of the all-domains (AD) configuration, where the EBA is only

0.01GHz/s better. In particular, it has a 12.26% lower L2-norm than EBA and 15.67% lower than VBA. The Average application execution progress ($AV-Wl_p$) with the FCA is 3.34% higher than EBA and 3.77% higher than VBA, while the minimum application execution progress ($MIN-Wl_p$) is 17.50% and 14.39% better in FCA than EBA and VBA respectively.

Both the EBA and VBA algorithms improve their target compliance and application progress results as the number of domains grows, managing to reach the performance of the FCA in the AD configuration, confirming the capability of EBA in absence of voltage coupling. In the MULTI-WL workload test, the VBA has a better $AV-Wl_p$ result but does worse in target compliance and $MIN-Wl_p$. This is because VBA tends to homogenize frequencies among cores, independently of the workload, favoring the less power-hungry (in this case the cores running the mix of int and floating point workload at 2.7GHz). Instead, FCA and EBA algorithms favor the cores running the more demanding workload, which, at the same conditions (F , V , T), consume more power and generates more heat than other workloads [9, 52].

Regarding results variation, Figure 9 shows that FCA has the smaller deviation in the average and minimum execution progress, while all three algorithms have a bimodal target compliance distribution. This is caused by the MULTI-WL tests. These tests have the majority of cores in a low-power state, thus leaving thermal and power room for the other cores to execute faster, and generating a set of tests in which all three algorithms perform better than the other types of workload (MAX-WL and CLOUD-WL). Figures (11, 12, 13) also shows that the FCA is the most consistent regarding target compliance and execution progress across different domains, models, and initial conditions respectively, with the variation being mostly regarding different workload. This is confirmed by Figure 10 where the error bars of FCA are significantly smaller than the other algorithms.

Summary. The average results and their standard deviation to account for results' consistency are summarized in Table 3.

The FCA achieves the best target compliance and application progress while outperforming the other algorithms in thermal capping and being able to cap power comparably to the VBA. This means that not only the FCA applies the given HLC target better, but it also makes the applications run around 3% faster and with lesser thermal room. Additionally, the execution progress across PEs is more consistent, as the application with the minimum progress is closer to the average value with respect to the other algorithms. The FCA has also the most consistent behavior across different cooling configurations, number of domains, and initial conditions, as the error bars in Figure 10 clearly show.

EBA total exceeding time distribution spread confirms that this algorithm doesn't have a consistent behavior when non-idealities are introduced into the model, as found also in Figure 7. Specifically, EBA and VBA encounter significant challenges with larger domains, highlighting the difficulties in managing voltage coupling. Conversely, the FCA's iterative root-finding method demonstrates better precision in achieving the desired operating point.

Power capping results suggest that it is difficult to effectively cap power and there is a trade-off between exceeding time and exceeding value, as all three algorithms had opposite behavior on those metrics .

7.2 HIL evaluation

While the focus of this work is to provide a cross-evaluation between several algorithms and thermal models of a modern computing HPC processor (Figure 3), we complete the assessment by verifying the feasibility of the proposed algorithms on the HIL framework proposed in [39] to exploit real hardware interfaces.

Alg			FCA	EBA	VBA
Temperature	Exc. Max Value [$^{\circ}\text{C}$]	Max	5.83	20.02	38.01
		Av	0.58	1.94	5.11
		SD	2.83	3.36	5.28
	Exceeded Time [%]	Av	0.38	18.82	4.01
		SD	0.73	12.10	7.09
	Average Temp [$^{\circ}\text{C}$]		67.87	71.96	69.23
Power	Exc. Av Value [%]	Av	13.98	3.74	15.14
		SD	10.65	3.45	25.16
	Exceeded Time [%]	Av	3.25	32.84	1.76
		SD	8.22	13.27	2.11
Execution	2-norm [GHz/s]	Av	0.80	0.91	0.95
		SD	0.35	0.45	0.45
	AV- Wl_p [%]	Av	78.10	75.57	75.26
		SD	8.55	12.01	12.58
	MIN- Wl_p [%]	Av	72.98	62.11	63.80
		SD	10.86	18.68	17.16

Table 3. Summary of the average results of the battery of tests. Standard deviation is included to show the constancy of each algorithm among different tests, models, and configurations. In green the best result, in red the worst.

On the controller side, the adoption of an FPGA emulation hosting the actual power controller HW offers significant advantages, such as the ability to test control firmware developed during the software- and model-in-the-loop (Software in the Loop (SIL) and MIL, respectively) design phases on actual hardware controllers with the assurance of cycle-accurate simulation. This precision is essential for achieving detailed hardware observability and controllability [51], as well as ensuring a one-to-one correspondence between the register transfer level (RTL) source and its FPGA implementation in terms of clock cycles, exemplified by direct FAME systems [55].

The framework consists of a Xilinx Zynq Ultrascale+ FPGA-SoC that hosts both the thermal and power model for the system under control and the LLC. We run the FCA on top of the LLC with a hyperperiod of 500us.

On a single core running at 50MHz, a basic version of the LLC firmware controls a chiplet of 24 cores, occupying up to the 78.18% of the hyperperiod. The FCA algorithm occupies on average the 96.3% of those cycles. In particular, of those cycles, the initialization contributes to the 9.57%, the Fuzzy-inspired thermal capping to the 10.54%, the Conv2P and the power capping to the 29.09%, and the iterative Conv2F part to the 48.79%. The remaining 2.02% comes from the sum of other initialization and management parts.

From this data, we can observe how the control algorithm is costly, and that the parts relative to the power conversion (Conv2P and Conv2F) take more than 50% of the whole algorithm. For this reason, the LLC HW proposed in [39] includes a cluster accelerator comprising 8 32-bit RISC-V cores to accelerate the more demanding fragments of the algorithm, such as the root-finding bisection

method.

The memory footprint of this basic firmware version is 120KiB. This includes initialized data, instructions, and uninitialized data. We observe that this fits the memory requirements of LLCs, typically below 1MiB.

This experiment demonstrates that (i) HW constraints such as execution speed and memory footprint are crucial to fully validate the developed controller algorithm; (ii) increasingly complex algorithms can be successfully implemented in HW, but timing constraints (control period) should be adjusted accordingly to fit the computation; (iii) more computationally-demanding algorithms, such as model-based or ML-driven, would be challenging to implement on actual LLCs, hence demanding for the integration of general-purpose or domain-specific HW acceleration techniques to meet the imposed constraints [39].

8 CONCLUSION

This paper presents the modeling of an HPC chip that takes into account the non-linearities and non-idealities of a real system and analyzes the control problem and challenges that arise from the realistic scenario. We first extend the work presented in Bambini et al. [7] to satisfy the introduced constraints and coupling in the thermal and power model. Then, we propose a novel control design inspired by the fuzzy control theory for thermal capping. Finally, we implement an iterative solving method to solve the coupling constraint instead of moving average filters.

The presented control design outperforms the PID algorithms in thermal capping by reducing the average exceeded temperature by up to 5× while being able to deliver more than 3% better application execution performance. Across all tests and configurations executed within a MIL framework, it proves to be the algorithm achieving the most stable results. The larger performance gains compared to the other algorithms happen in the scenario with larger domains (1-D and 4-D) and the RACK model, where the thermal coupling is more complex. The analysis of the results suggests that the proposed design is beneficial in addressing the control challenges presented in Section 4.2. The original design from [7] is still competitive in the A-D configuration, being the more compliant with the assumptions made in the original work.

While the fuzzy-inspired design solves the thermal capping problem, the power distribution block is still overly complex: it is model-based, it assumes to have a measure of the workload ω , and it does not fully take into account the output coupling. Future works in this direction should investigate a simpler and more comprehensive power distribution block. Model predictive solutions could be considered, provided their efficient deployment on resource-constrained LLCs.

9 ACKNOWLEDGMENTS

The study has been conducted in the context of EU H2020-JTI-EuroHPC-2019-1 project REGALE (g.n. 956560), EuroHPC EU PILOT project (g.a. 101034126), EU Pilot for exascale EuroHPC EUPLEX (g. a. 101033975), and European Processor Initiative (EPI) SGA2 (g.a. 101036168).

REFERENCES

- [1] AMD 2013. *BIOS and Kernel Developer's Guide (BKDG) for AMD Family 15h Models 00h-0Fh Processors*. AMD.
- [2] ARM 2022. *Arm® Neoverse™ V2 Core Technical Reference Manual 5.5.1*. ARM. <https://developer.arm.com/documentation/102375/latest/>
- [3] ARM 2023. *Power Control System Architecture - DEN0050*. ARM. <https://developer.arm.com/documentation/den0050/latest/>
- [4] Arm. 2023. SCP-firmware - version 2.13. <https://github.com/Arm-software/SCP-firmware>.
- [5] Eberhard Baer, Alex Burenkov, Peter Evanschitzky, and Juergen Lorenz. 2016. Simulation of process variations in FinFET transistor patterning. In *2016 International Conference on Simulation of Semiconductor Processes and Devices (SISPAD)*. 299–302. <https://doi.org/10.1109/SISPAD.2016.7605206>

- [6] Giovanni Bambini, Robert Balas, Christian Conficoni, Andrea Tilli, Luca Benini, Simone Benatti, and Andrea Bartolini. 2020. An Open-Source Scalable Thermal and Power Controller for HPC Processors. In *2020 IEEE 38th International Conference on Computer Design (ICCD)*. 364–367. <https://doi.org/10.1109/ICCD50377.2020.00067>
- [7] Giovanni Bambini, Christian Conficoni, Andrea Tilli, Luca Benini, and Andrea Bartolini. 2022. Modeling the Thermal and Power Control Subsystem in HPC Processors. In *2022 IEEE Conference on Control Technology and Applications (CCTA)*. 397–402. <https://doi.org/10.1109/CCTA49430.2022.9966082>
- [8] A. Bartolini, M. Cacciari, A. Tilli, and L. Benini. 2013. Thermal and Energy Management of High-Performance Multicores: Distributed and Self-Calibrating Model-Predictive Controller. *IEEE Transactions on Parallel and Distributed Systems* 24, 1 (2013), 170–183.
- [9] Andrea Bartolini and Davide Rossi. 2019. Advances in power management of many-core processors. *Many-Core Computing: Hardware and Software* (2019), 191.
- [10] Anton Belyazov, Rajkumar Buyya, Young Choon Lee, and Albert Zomaya. 2011. Chapter 3 - A Taxonomy and Survey of Energy-Efficient Data Centers and Cloud Computing Systems. *Advances in Computers*, Vol. 82. Elsevier, 47–111. <https://doi.org/10.1016/B978-0-12-385512-1.00003-7>
- [11] Abdelhalim Bendali and Yves Audet. 2007. A 1-V CMOS Current Reference With Temperature and Process Compensation. *IEEE Transactions on Circuits and Systems I: Regular Papers* 54, 7 (2007), 1424–1429. <https://doi.org/10.1109/TCSI.2007.900176>
- [12] F. Beneventi, A. Bartolini, A. Tilli, and L. Benini. 2014. An Effective Gray-Box Identification Procedure for Multicore Thermal Modeling. *IEEE Trans. Comput.* 63, 5 (2014), 1097–1110.
- [13] Ganapati Bhat, Gaurav Singla, Ali K. Unver, and Umit Ogras. 2017. Algorithmic Optimization of Thermal and Power Management for Heterogeneous Mobile Platforms. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems* (2017). <https://doi.org/10.1109/TVLSI.2017.2770163>
- [14] Thomas Burd, Noah Beck, Sean White, Milam Paraschou, Nathan Kalyanasundharam, Gregg Donley, Alan Smith, Larry Hewitt, and Samuel Naffziger. 2019. “Zeppelin”: An SoC for Multichip Architectures. *IEEE Journal of Solid-State Circuits* 54, 1 (2019), 133–143. <https://doi.org/10.1109/JSSC.2018.2873584>
- [15] Richard L Burden, J Douglas Faires, and Annette M Burden. 2015. *Numerical analysis*. Cengage learning.
- [16] Chang-Chih Chen and Linda Milor. 2015. Microprocessor Aging Analysis and Reliability Modeling Due to Back-End Wearout Mechanisms. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems* 23, 10 (2015), 2065–2076. <https://doi.org/10.1109/TVLSI.2014.2357756>
- [17] Hsiang-Yun Cheng, Jia Zhan, Jishen Zhao, Yuan Xie, Jack Sampson, and Mary Jane Irwin. 2015. Core vs. uncore: The heart of darkness. In *2015 52nd ACM/EDAC/IEEE Design Automation Conference (DAC)*. 1–6. <https://doi.org/10.1145/2744769.2647916>
- [18] Yingnan Cui, Wei Zhang, and Bingsheng He. 2017. A Variation-Aware Adaptive Fuzzy Control System for Thermal Management of Microprocessors. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems* 25, 2 (2017), 683–695. <https://doi.org/10.1109/TVLSI.2016.2596338>
- [19] Pawel Czarul, Jerzy Proficz, and Adam Krzywaniak. 2019. Energy-Aware High-Performance Computing: Survey of State-of-the-Art Tools, Techniques, and Environments. *Scientific Programming* 2019 (04 2019), 1–19. <https://doi.org/10.1155/2019/8348791>
- [20] Shidhartha Das, Paul Whatmough, and David Bull. 2015. Modeling and characterization of the system-level Power Delivery Network for a dual-core ARM Cortex-A57 cluster in 28nm CMOS. In *2015 IEEE/ACM International Symposium on Low Power Electronics and Design (ISLPED)*. 146–151. <https://doi.org/10.1109/ISLPED.2015.7273505>
- [21] Anant Deval, Avinash Ananthkrishnan, and Craig Forbell. 2015. Power management on 14 nm Intel® Core M processor. *2015 IEEE Symposium in Low-Power and High-Speed Chips (COOL CHIPS XVIII)* (2015), 1–3. <https://api.semanticscholar.org/CorpusID:37333321>
- [22] Roberto Diversi, Andrea Tilli, Andrea Bartolini, Francesco Beneventi, and Luca Benini. 2014. Bias-Compensated Least Squares Identification of Distributed Thermal Models for Many-Core Systems-on-Chip. *IEEE Transactions on Circuits and Systems I: Regular Papers* 61, 9 (2014), 2663–2676. <https://doi.org/10.1109/TCSI.2014.2312495>
- [23] Lawrence C. Evans. 2010. *Partial differential equations*. American Mathematical Society, Providence, R.I.
- [24] Gene Franklin, J.D. Powell, and M.L. Workman. 2022. *Digital Control of Dynamic Systems-Third Edition*.
- [25] Bolin Gao and Laca Pavel. 2018. On the Properties of the Softmax Function with Application in Game Theory and Reinforcement Learning. *arXiv:1704.00805 [math.OC]*
- [26] Programmer Guide. 2022. *Intel 64 and IA-32 Architectures Software Developer’s Manual*. Intel. <https://software.intel.com/content/www/us/en/develop/articles/intel-sdm.html>
- [27] Daniel Hackenberg, Robert Schöne, Thomas Ilsche, Daniel Molka, Joseph Schuchart, and Robin Geyer. 2015. An Energy Efficiency Feature Survey of the Intel Haswell Processor. In *2015 IEEE International Parallel and Distributed Processing Symposium Workshop*. 896–904. <https://doi.org/10.1109/IPDPSW.2015.70>
- [28] Vinay Hanumaiah, Sarma Vrudhula, and Karam S. Chatha. 2011. Performance Optimal Online DVFS and Task

- Migration Techniques for Thermally Constrained Multi-Core Processors. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 30, 11 (2011), 1677–1690. <https://doi.org/10.1109/TCAD.2011.2161308>
- [29] IBM. 2022. OpenPower OCC. <https://github.com/open-power/occ>.
- [30] Seri Lee and Kevin P. Moran. 1996. Constriction/spreading resistance model for electronics packaging. <https://api.semanticscholar.org/CorpusID:28843083>
- [31] Alberto Leva, Federico Terraneo, Irene Giacomello, and William Fornaciari. 2018. Event-Based Power/Performance-Aware Thermal Management for High-Density Microprocessors. *IEEE Transactions on Control Systems Technology* 26, 2 (2018), 535–550. <https://doi.org/10.1109/TCST.2017.2675841>
- [32] Zhifeng Liu and Hong Zhu. 2010. A survey of the research on power management techniques for high-performance systems. *Softw., Pract. Exper.* 40 (10 2010). <https://doi.org/10.1002/spe.v40:11>
- [33] Google LLC. U.S. Patent US8402290B2, Dec. 2020. Power management for multiple processor cores.
- [34] Abhinandan Majumdar, Leonardo Piga, Indrani Paul, Joseph L. Greathouse, Wei Huang, and David H. Albonesi. 2017. Dynamic GPGPU Power Management Using Adaptive Model Predictive Control. In *2017 IEEE International Symposium on High Performance Computer Architecture (HPCA)*. 613–624. <https://doi.org/10.1109/HPCA.2017.34>
- [35] Sumit K. Mandal, Ganapati Bhat, Janardhan Rao Doppa, Partha Pratim Pande, and Umit Y. Ogras. 2020. An Energy-aware Online Learning Framework for Resource Management in Heterogeneous Platforms. *ACM Trans. Des. Autom. Electron. Syst.* 25, 3, Article 28 (may 2020), 26 pages. <https://doi.org/10.1145/3386359>
- [36] Kasra Moazzemi, Biswadip Maity, Saehanseul Yi, Amir M. Rahmani, and Nikil Dutt. 2019. HESSLE-FREE: Heterogeneous Systems Leveraging Fuzzy Control for Runtime Resource Management. *ACM Trans. Embed. Comput. Syst.* 18, 5s, Article 74 (oct 2019), 19 pages. <https://doi.org/10.1145/3358203>
- [37] Samuel Naffziger, Noah Beck, Thomas Burd, Kevin Lepak, Gabriel H. Loh, Mahesh Subramony, and Sean White. 2021. Pioneering Chiplet Technology and Design for the AMD EPYC™ and Ryzen™ Processor Families : Industrial Product. In *2021 ACM/IEEE 48th Annual International Symposium on Computer Architecture (ISCA)*. 57–70. <https://doi.org/10.1109/ISCA52012.2021.00014>
- [38] Hung T Nguyen and Michio Sugeno. 2012. *Fuzzy systems: modeling and control*. Vol. 2. Springer Science & Business Media.
- [39] Alessandro Ottaviano, Robert Balas, Giovanni Bambini, Antonio Del Vecchio, Maicol Ciani, Davide Rossi, Luca Benini, and Andrea Bartolini. 2024. ControlPULP: A RISC-V On-Chip Parallel Power Controller for Many-Core HPC Processors with FPGA-Based Hardware-In-The-Loop Power and Thermal Emulation. *International Journal of Parallel Programming* (26 Feb 2024). <https://doi.org/10.1007/s10766-024-00761-4>
- [40] G. Paci, P. Marchal, F. Poletti, and L. Benini. 2006. Exploring “temperature-aware” design in low-power MPSoCs. In *Proceedings of the Design Automation & Test in Europe Conference*, Vol. 1. 1–6. <https://doi.org/10.1109/DATE.2006.243741>
- [41] Santiago Pagani, P. D. Sai Manoj, Axel Jantsch, and Jörg Henkel. 2020. Machine Learning for Power, Energy, and Thermal Management on Multicore Processors: A Survey. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 39, 1 (2020), 101–116. <https://doi.org/10.1109/TCAD.2018.2878168>
- [42] Jaehyun Park, Donghwa Shin, Naehyuck Chang, and Massoud Pedram. 2010. Accurate modeling and calculation of delay and energy overheads of dynamic voltage scaling in modern high-performance microprocessors. In *2010 ACM/IEEE International Symposium on Low-Power Electronics and Design (ISLPED)*. 419–424. <https://doi.org/10.1145/1840845.1840938>
- [43] Jan M Rabaey. 1999. *Digital integrated circuits a design perspective*.
- [44] Martin Rapp, Mohammed Bakr Sikal, Heba Khdr, and Jörg Henkel. 2021. SmartBoost: Lightweight ML-Driven Boosting for Thermally-Constrained Many-Core Processors. In *2021 58th ACM/IEEE Design Automation Conference (DAC)*. 265–270. <https://doi.org/10.1109/DAC18074.2021.9586287>
- [45] RISC-V 2017. *The RISC-V Instruction Set Manual Volume II: Privileged Architecture*. RISC-V. <https://riscv.org/technical/specifications/>
- [46] Chiara Irma Riva. 2021. A numerical tool for the analytical solution of temperature rise and thermal spreading resistance for power electronics. (2021).
- [47] Todd Rosedahl, Martha Broyles, Charles Lefurgy, Bjorn Christensen, and Wu Feng. 2017. Power/Performance Controlling Techniques in OpenPOWER. In *High Performance Computing*, Julian M. Kunkel, Rio Yokota, Michela Taufer, and John Shalf (Eds.). Springer International Publishing, Cham, 275–289.
- [48] Davide Rossi, Antonio Pullini, Igor Loi, Michael Gautschi, Frank K. Gürkaynak, Andrea Bartolini, Philippe Flatresse, and Luca Benini. 2016. A 60 GOPS/W, -1.8V to 0.9V body bias ULP cluster in 28nm UTBB FD-SOI technology. *Solid-State Electronics* 117 (2016), 170–184. <https://doi.org/10.1016/j.sse.2015.11.015> PLANAR FULLY-DEPLETED SOI TECHNOLOGY.
- [49] Efraim Rotem, Ran Ginosar, Avi Mendelson, and Uri C. Weiser. 2013. Power and thermal constraints of modern system-on-a-chip computer. In *19th International Workshop on Thermal Investigations of ICs and Systems (THERMINIC)*. 141–146. <https://doi.org/10.1109/THERMINIC.2013.6675226>

- [50] Efraim Rotem, Alon Naveh, Avinash Ananthakrishnan, Eliezer Weissmann, and Doron Rajwan. 2012. Power-Management Architecture of the Intel Microarchitecture Code-Named Sandy Bridge. *IEEE Micro* 32, 2 (2012), 20–27. <https://doi.org/10.1109/MM.2012.12>
- [51] Martin Schlager, Roman Obermaisser, and Wilfried Elmenreich. 2007. A Framework for Hardware-in-the-Loop Testing of an Integrated Architecture. *Lecture Notes in Computer Science - LNCS* 4761, 159–170. https://doi.org/10.1007/978-3-540-75664-4_16
- [52] Robert Schöne, Thomas Ilsche, Mario Bielert, Andreas Gocht, and Daniel Hackenberg. 2019. Energy Efficiency Features of the Intel Skylake-SP Processor and Their Impact on Performance. *arXiv:1905.12468 [cs]* (May 2019). <http://arxiv.org/abs/1905.12468> arXiv: 1905.12468.
- [53] Kevin Skadron, Mircea R. Stan, Wei Huang, Sivakumar Velusamy, Karthik Sankaranarayanan, and David Tarjan. 2003. Temperature-Aware Microarchitecture. In *Proceedings of the 30th Annual International Symposium on Computer Architecture* (San Diego, California) (*ISCA '03*). Association for Computing Machinery, New York, NY, USA, 2–13. <https://doi.org/10.1145/859618.859620>
- [54] Hameedah Sultan, Anjali Chauhan, and Smruti R. Sarangi. 2019. A Survey of Chip-Level Thermal Simulators. *ACM Comput. Surv.* 52, 2, Article 42 (apr 2019), 35 pages. <https://doi.org/10.1145/3309544>
- [55] Zhangxi Tan, Andrew Waterman, Henry Cook, Sarah Bird, Krste Asanović, and David Patterson. 2010. A Case for FAME: FPGA Architecture Model Execution. In *Proceedings of the 37th Annual International Symposium on Computer Architecture* (Saint-Malo, France) (*ISCA '10*). Association for Computing Machinery, New York, NY, USA, 290–301. <https://doi.org/10.1145/1815961.1815999>
- [56] Andrea Tilli, Emanuele Garone, Christian Conficoni, Matteo Cacciari, Alessandro Bosso, and Andrea Bartolini. 2022. A two-layer distributed MPC approach to thermal control of Multiprocessor Systems-on-Chip. *Control Engineering Practice* 122 (5 2022). <https://doi.org/10.1016/j.conengprac.2022.105099>
- [57] Ankush Varma, Bill Bowhill, Jason Crop, Corey Gough, Brian Griffith, Dan Kingsley, and Krishna Sistla. 2015. Power management in the Intel Xeon E5 v3. In *2015 IEEE/ACM International Symposium on Low Power Electronics and Design (ISLPED)*. 371–376. <https://doi.org/10.1109/ISLPED.2015.7273542>
- [58] A. Vassighi and M. Sachdev. 2006. Thermal runaway in integrated circuits. *IEEE Transactions on Device and Materials Reliability* 6, 2 (2006), 300–305. <https://doi.org/10.1109/TDMR.2006.876577>
- [59] Wei-Bin Yang, Yu-Yao Lin, and Yu-Lung Lo. 2014. Analysis and design considerations of static CMOS logics under process, voltage and temperature variation in 90nm CMOS process. In *2014 International Conference on Information Science, Electronics and Electrical Engineering*, Vol. 3. 1653–1656. <https://doi.org/10.1109/InfoSEE.2014.6946202>
- [60] Huazhe Zhang and Henry Hoffmann. 2016. Maximizing Performance Under a Power Cap: A Comparison of Hardware, Software, and Hybrid Techniques. 51, 4 (mar 2016), 545–559. <https://doi.org/10.1145/2954679.2872375>