


REliable Power and time-ConstraInts-aware Predictive management of heterogeneous
Exascale systems



WP3 Predictive Reliability and QoS Enforcing Methodologies

D3.5 RECIPE Thermal Simulation Tools

<http://www.recipe-project.eu>



This project has received funding from the European Union's Horizon
2020 research and innovation programme under grant agreement No
801137

Grant Agreement No.: 801137

Deliverable: D3.5 RECIPE Thermal Simulation Tools

Project Start Date: 01/05/2018

Duration: 36 months

Coordinator: *Politecnico di Milano, Italy*

Deliverable No:	D3.5
WP No:	3
WP Leader:	R. Canal
Due date:	30/04/2020
Delivery date:	07/05/2020

Dissemination Level:

PU	Public Use	X
PP	Restricted to other programme participants (including the Commission Services)	
RE	Restricted to a group specified by the consortium (including the Commission Services)	
CO	Confidential, only for members of the consortium (including the Commission Services)	

DOCUMENT SUMMARY INFORMATION

Project title:	REliable Power and time-ConstraInts-aware Predictive management of heterogeneous Exascale systems
Short project name:	RECIPE
Project No:	801137
Call Identifier:	H2020-FETHPC-2017
Thematic Priority:	Future and Emerging Technologies
Type of Action:	Research and Innovation Action
Start date of the project:	01/05/2018
Duration of the project:	36 months
Project website:	http://www.recipe-project.eu

D3.5 RECIPE Thermal Simulation Tools

Work Package:	WP3 Predictive Reliability and QoS Enforcing Methodologies
Deliverable number:	D3.5
Deliverable title:	RECIPE Thermal Simulation Tools
Due date:	30/04/2020
Actual submission date:	07/05/2020
Editor:	M. Zapater
Authors:	M. Zapater, D. Atienza, W. Piatek, S. Ciesielski, W. Szeliga, A. Oleksiak, M. Kulczewski
Dissemination Level:	PU
No. pages:	21
Authorized (date):	07/05/2020
Responsible person:	W. Fornaciari
Status:	Submitted

Revision history:

Version	Date	Author	Comment
v.0.1	13/04/2020		Outline and contributions identified

Quality Control:

	Who	Date
Checked by internal reviewer	M. Zapater	04/05/2020
Checked by WP Leader	R. Canal	06/05/2020
Checked by Project Technical Manager	G. Agosta	07/05/2020
Checked by Project Coordinator	W. Fornaciari	07/05/2020

COPYRIGHT

©Copyright by the **RECIPE** consortium, 2018-2020.

This document contains material, which is the copyright of RECIPE consortium members and the European Commission, and may not be reproduced or copied without permission, except as mandated by the European Commission Grant Agreement no. 801137 for reviewing and dissemination purposes.

ACKNOWLEDGEMENTS

RECIPE is a project that has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement No 801137. Please see <http://www.recipe-project.eu> for more information.

The partners in the project are Universitat Politècnica de València (UPV), Centro Regionale Information Communication Technology srl (CeRICT), École Polytechnique Fédérale de Lausanne (EPFL), Barcelona Supercomputing Center (BSC), Poznan Supercomputing and Networking Center (PSNC), IBT Solutions S.r.l. (IBTS), Centre Hospitalier Universitaire Vaudois (CHUV). The content of this document is the result of extensive discussions within the RECIPE ©Consortium as a whole.

DISCLAIMER

The content of the publication herein is the sole responsibility of the publishers and it does not necessarily represent the views expressed by the European Commission or its services. The information contained in this document is provided by the copyright holders "as is" and any express or implied warranties, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose are disclaimed. In no event shall the members of the RECIPE collaboration, including the copyright holders, or the European Commission be liable for any direct, indirect, incidental, special, exemplary, or consequential damages (including, but not limited to, procurement of substitute goods or services; loss of use, data, or profits; or business interruption) however caused and on any theory of liability, whether in contract, strict liability, or tort (including negligence or otherwise) arising in any way out of the use of the information contained in this document, even if advised of the possibility of such damage.

Contents

1	Introduction	6
2	Summary of the Thermal Modelling Methodology	7
2.1	System-on-Chip modelling	7
2.2	Server/farm modelling	9
3	Thermal Modelling Tools	18
3.1	DCworms	18
4	Summary	19

1 Introduction

High-Performance Computing (HPC) systems have become ubiquitous and are no longer concentrated in supercomputing facilities and data centers. While these facilities still exist and grow, a plethora of HPC systems building on large multicores and accelerators (e.g. GPUs, FPGA-based) are nowadays deployed for a variety of applications of interest, not only for large enterprises and public institutions, but also for small and medium enterprises as well as small public and private bodies.

The proliferation of HPC systems and applications in new domains has led to new requirements related to non-functional requirements (time, power, reliability, temperature, etc) and the implementation of platforms to satisfy them [1, 7, 6, 10, 13]. In this deliverable, we provide a description of the thermal simulation tools developed in this project. Temperature prediction (or thermal behaviour modelling) is key to provide good reliability estimates, but it is also important to leverage cooling capabilities at the board, rack and room levels to optimize its efficiency. By utilising the power of CFD simulations combined with experiments conducted on real hardware we create corresponding power and thermal models that can be later on applied to study the energy-efficiency of the whole data centre. Our simulation tools are based on two (interconnected) levels:

- Thermal modelling techniques to estimate chip temperature. Our work in T3.3 provides means to simulate and model chip temperature accurately with the goal of proactively enhancing chip reliability.
- Thermal modelling techniques to estimate board, rack and room temperature and cooling facilities efficiency. Our work in T3.3 provides means to simulate and model whole system temperature behaviour with the goal of proactively enhancing cooling efficiency.

The rest of the deliverable provides details on the usage and capabilities of the software developed. The specific underlying techniques implemented are described in the previous deliverable D3.2.

The tools developed are available at:

- DCworms: <https://git.man.poznan.pl/stash/projects/WORMS/repos/dcworms>

2 Summary of the Thermal Modelling Methodology

2.1 System-on-Chip modelling

Thermal modelling within RECIPE is required to first enable the development of thermal-aware policies but also, and more importantly, to enable the evaluation of thermal effects on the long-term reliability of servers. Within WP3, thermal models of processors and servers are developed, that are then used in WP2 for the development of thermal and reliability aware management policies at the server level.

In particular, thermal stress and spatial gradients are proven to have a negative effect on the long-term reliability of Multi-Processors System on Chip (MPSoCs) [2], impacting their FIT rate [9]. Temporal Temperature Gradient (TTG) can be defined as the rate of temperature changes over time. For a given time, the rate of the temperature changes from one point to another indicates the spatial temperature gradient (STG). Both STG and TTG pose a critical impact on the system lifetime reliability, but STG is mostly affected by the power and thermal management techniques applied at the overall MPSoC system, i.e., the allocation and specific setup of all cores in the system need to be taken into consideration. In contrast, TTG is mostly affected by the core frequency and the workload running in each specific core.

Thermal cycling is the phenomena which takes place when the temperature rises up (or drops down) and goes back to the initial value (which can be defined as a thermal cycle) frequently [16]. MTTF reduction due to thermal cycling occurs due to the mismatch on the expansion coefficient between the layers of the chip, which results in thermo-mechanical stresses. Thermal cycling (TC) tends to reduce the whole system MTTF as the number of cycles or amplitudes increases. Large amplitudes are normally induced due to improper task scheduling on a single core. Number of thermal cycles increases especially by the power management techniques which frequently turn cores on and off [3].

All in all, in order to reliably estimate the MTTF of a system, we need ways of modelling and simulating temperature in a spatio-temporal way. In this sense, simulators such as the 3D-ICE tool developed as part of the previous work of EPFL can help in accurately modelling these effects [14]. However, in order for this tool to work efficiently and accurately within the framework of RECIPE, there is a need to incorporate the following enhancements:

- Enabling the simulation of arbitrary state-of-the-art cooling mechanisms, such as the ones found in current servers. In particular, within RECIPE we need to enable the simulation of both natural convection mechanisms (i.e., heatsinks), and forced convection cooling (i.e., heatsinks plus fans).
- Proposal of a methodology that will allow us to assess the impact of the main control knobs related to temperature in today's servers, which range from workload allocation, DVFS setting and fan speed control policies. This methodology needs to exploit the capabilities of our simulation tool.
- Proposal of a methodology to adequately link the thermal aspects to the reliability of the system. For this purpose, we will use the MTTF and reliability models proposed by T3.1, which will be incorporated into the policies developed by EPFL in both T3.5 in WP3 and T2.3 in WP2.

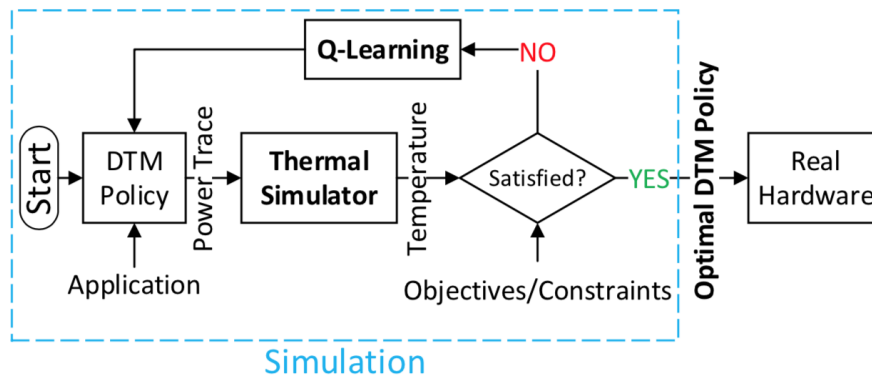


Figure 1: Thermal modelling methodology to enable DTM and reliability management

Based on the previous work released in v2.2.6 of 3D-ICE (project background), within this task of RECIPE, we have developed both the natural convection (heatspreader plus heatsink) and the forced convection (heatsink plus fan) models, using a real Thermal Test Chip (TTV) , for real chips and cooling devices, using real traces. This contributions are released in v2.2.7 of 3D-ICE. We plan to release v2.2.8 in the near future, incorporating both the natural and forced convection models developed. The details of the pluggable heatsink and its integration with 3D-ICE will be described as part of Deliverable D3.5.

Furthermore, to create accurate models of the system, we need ways of validating them against real devices. For this purpose, within the RECIPE project we have also created a platform which comprises a real test chip [15] for accurate thermal characterization. In particular, this platform is based on a Thermal Test Chip (TTC), an integrated circuit containing an array of power dissipating elements and an array of temperature sensors. Our thermal platform is capable of applying a generic power dissipation pattern to the thermal test chip and measuring the corresponding temperature map, at a rate up to 1 kHz. This capability allows to measure the temperature map of an integrated circuit subject to reference power dissipation maps, and thus to design and validate thermal models. In our particular case, the chip is organised as a 4 by 4 array of individual cells, each capable of temperature sensing and power generation through a resistive element. The heating element in each cell is capable of dissipating up to 12 W. This allows to simulate the equivalent of a 16-core chip with a total power dissipation of 192 W.

The methodology for envision for enabling DTM and reliability management is the one depicted in Figure 1. The simulator will enable us perform an exploration of the impact of aggressive DTM management policies which may not only cause performance degradation and additional power consumption, but more importantly, jeopardizes lifetime reliability of the whole system. In fact, one of the main reasons that makes researchers reluctant to consider fan speed control as a key DTM approach is the lack of a transient thermal simulator for MPSoCs with proper integration of fans. The incorporation of fan models in 3D-ICE, and the use of this methodology provides a comprehensive framework for exploring thermal effects of DTM policies in a safe way.

Workload allocation, DVFS and fan speed altogether drastically increase the number of runtime design parameters to be decided by a DTM and reliability-aware policy, which leads to additional challenges to find the best values for optimal behavior of the whole system. Our methodology enables the exploration of this design space in an automated way. In particular, we envision the

use of Reinforcement Learning (RL) techniques in WP2. The RL agents can explore the design space using 3D-ICE and combine that with the impact of reliability. Initial experiments [15], prove the feasibility of our approach.

2.2 Server/farm modelling

Apart from modelling the system on a chip level, there is a need to evaluate its impact not only on the server to which it belongs but also refer it to a large scale computing infrastructure while taking into account the corresponding facilities like cooling equipment. Thus, the main idea is to create power and thermal models to describe the behaviour of a single server and its subcomponents, and then transpose them to describe the behaviour of the whole system (rack, server or data centre) consisting of the given server. On one hand, these models need to reflect the operation of real hardware with the necessary accuracy and, on the other hand, enable their application in the evaluation of the large-scale system in an efficient manner. The models will be based on data obtained via real experiments, CFD simulations and supported with the data coming from the manufacturer's specifications. They will also rely on the output of 3D-ICE thermal simulation tool. As the tool considers detailed specification of SoCs, it could provide more coarse-grained thermal characteristics as well.

In order to conduct experiments and gather data necessary for creating thermal characteristics of the nodes, we adopted a benchmarking tool developed in PSNC. It allows stressing the nodes with a given application and collecting their characteristics. For the purpose of our studies, we adjusted this tool to gather also data from sensors available within the evaluated SuperServer system.

The parameters collected by the tool include power and temperature of each of the evaluated module. Additionally, using SuperServer interface (SuperDoctor 5), the tool tracks speed of fans and internal temperatures of the server. Finally, temperature of the air in server room is recorded for each experiment. The main goal of the performed experiments was to measure changes in power and temperature for various load levels. Each test consisted of two phases: 1) stressing the given node for 15 minutes and 2) putting the nodes in an idle state for 5 minutes. The former phase allowed observing the process of heating the node and reaching its stable state (for a given load), while the latter highlighted how the system is returning to its initial (idle) state (the process of cooling the node). One should note that the fans in the evaluated systems adjust their speed in an automatic manner, thus it is crucial to capture this habit to reflect their impact on the energy efficiency of the system.

In general, the testbed consists of two server systems. The first one contains two Intel Xeon Silver 4112 processors and an NVIDIA Tesla P100 GPU accelerator, while the second one comes with Stratix 10 FPGA module. For now, the experiments have been conducted on the former server. The latter server will be subject of the upcoming studies.

Figure 2 shows the location (red rectangles) of the computing resources and cooling fans inside the first server.

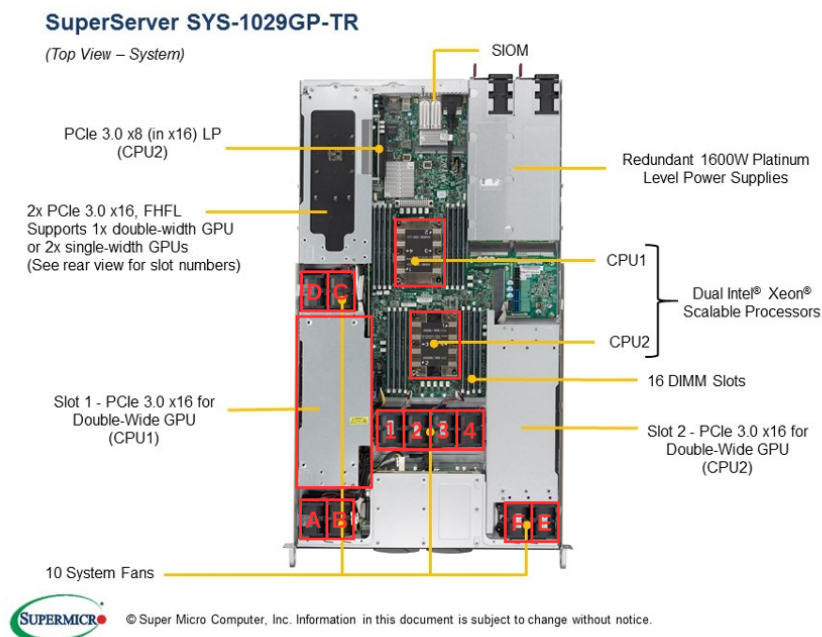


Figure 2: Arrangement of particular components inside the evaluated server

The following tests were carried out:

Test case	CPU load [%]	GPU load [%]
1	20	0
2	40	0
3	60	0
4	80	0
5	100	0
6	20	100
7	40	100
8	60	100
9	80	100
10	100	100
11	0	100

Table 1: A summary of conducted experiments

Figure 3 shows power consumption for both CPUs while changing the load from 20% to 100%. During these experiments, the GPU was in an idle state. These results correspond to test cases 1-5.

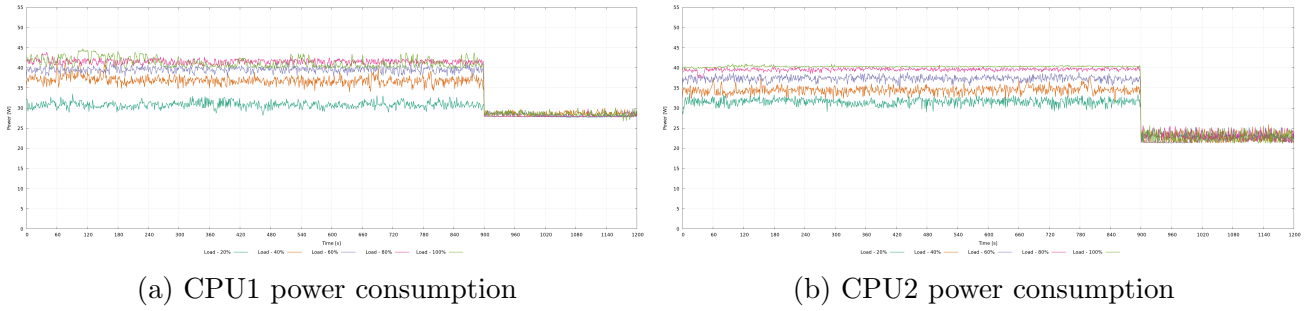


Figure 3: Power consumption for various load levels

Figure 4 compares power consumption of CPU1 and CPU2 for 100% load (test case 5).




Figure 4: Comparison of cpus' power for 100% load

One should note that the power characteristics for particular load levels differs between the CPUs. The "warmer" the CPU (as indicated in the figures below), the higher its power usage. This suggests that the power leakage phenomenon for the analysed CPUs cannot be negligible.

Figure 5 shows temperature levels for both CPUs while changing the load from 20% to 100%. These results also refer to test cases 1-5.




Figure 5: Temperature for various load levels

Figure 6 compares the temperatures of CPU1 and CPU2 for 100% load level (test case 5).




Figure 6: Comparison of cpus' temperature for 100% load

As expected, the temperature of CPU depends on its location inside the server. The closer the CPU to the inlet the lower the temperatures.

Figure 7 depicts the values of fans speeds when utilising the CPUs at 100% (test case 5). Fans with the same speed have been grouped. The location of fans is shown in Figure 2.




Figure 7: Changes in fans speed

One should note that fans change their speed gradually. When stressing the CPUs, only fans located in their area (Fans 1-4) take part in the cooling process.

Figure 8 shows power usage for 100% load level applied to GPU. During this experiment, the processors were idle. This scenario corresponds to test case 11.




Figure 8: GPU power consumption for 100% load

Next figure (Figure 9) presents the changes in temperature for test case 11.




Figure 9: GPU temperature for 100% load

Finally, Figure 10 presents how fans react when stressing the GPU. Fans with the same speed have been grouped. The location of fans is shown in Figure 2.




Figure 10: Changes in fans speed

Again, fans located close to the GPU are mainly responsible for cooling. Moreover, one can observe a gradual increase in their speed.

Apart from experiments, it is also intended to provide latter thermal models with results of numerical simulations. We plan to perform CFD analysis of the examined server for various scenarios, that samples different parameters such as fans rotation speed or load of server components. CFD analysis is a good complement to experiments, as collected data is volumetric and not limited to several sensors. Such simulations are required in situations where experimental examinations are not possible.

In order to run simulations, it is necessary to provide a virtual, CFD-suited model of the server. The model has been provided based on the available documentation. Exemplary visualisation of the model is presented in Figure 11.




Figure 11: Virtual model of SuperServer SYS-1029GP-TR; top cover is removed and air shroud over CPU heatsinks is made transparent, both for better model visibility

The model includes the most important components responsible for airflow and heat transfer formation inside the enclosure. The limitation encountered at this preliminary stage is the lack of information about internal structure of both the GPU accelerator and power supply boxes. Due to this fact they are now modelled as thermally passive devices with empty flow-through ducts inside, but in the future they can be improved with reasonable black box approach.

The preliminary simulation scenario described in this paragraph covers both CPU occupied with 100% load ($TDP = 85W$) and motherboard chipset heat generation rate also at the maximum level ($TDP = 15W$), while GPU and power supply devices are thermally passive. All the fans are rotating with the maximum speed. Server's inlet air temperature is equal to $20^{\circ}C$.

Simulation results provide both server's aerodynamics and heat transfer coverage. In Figure 12, the airflow direction inside the enclosure is presented using streamlines generated with respect to received velocity volumetric data. Coarse visual analysis has already revealed several spots, where the airflow can be optimised by sealing passages around fans in order to prevent backflows, reducing airflow and heat dissipation effectiveness, as indicated in Figure 13.




Figure 12: Server's aerodynamics analysis: streamlines describing air motion inside the enclosure




Figure 13: Backflows revealed in CFD simulation indicate the need of sealing gaps around fans

Besides kinematics of the air, the simulation covers heat transfer phenomenon as well. Exemplary visualization of temperature values at server's equipment surfaces is presented in Figure 14.




Figure 14: Temperature values visualized at server's equipment surfaces (temperature is given in Kelvin scale)

Based on results, for a case with maximum CPU load together with maximum rotation speed of all fans, temperature at CPU1 and CPU2 reaches 36.4°C and 43.0°C respectively. It is not intended to compare this particular simulation results with previously described experiments as the preliminary CFD simulation was performed for different conditions (fans rotation speed and inlet temperature).

As mentioned, the outcome of the CFD analysis results will be used to enrich the experimental data with the air distribution model. Moreover, it will support and complement the experimental studies with the analysis of edge cases and providing data difficult to measure via real tests or limited by the design of the evaluated system. Additionally, on a data center level, CFD analysis can be used to validate the scalability of the RECIPE solutions by exposing spatial information omitted by standard scheduling simulators. These may include the efficiency of the cooling medium, hot-spot detection and design validation.

Based on experimental and numerical data we will determine the proper, possibly minimum, set of factors that will allow describing the hardware behaviour from the power and thermal point of view. To this end, the duality between thermal and electrical phenomena, the law of energy conservation and the affinity laws for fans will be used. To characterise changes in CPU's temperature we will benefit from the following dependencies:

$$T_{cpu}(t + \Delta t) = T_{cpu}^{\infty}(t + \Delta t) + (T_{cpu}(t) - T_{cpu}^{\infty}(t + \Delta t))e^{-\frac{\Delta t}{R(t + \Delta t)C}} \quad (1)$$

with

$$T_{cpu}^{\infty}(t) = P_{cpu}(t)R(t) + T_{amb}(t) \quad (2)$$

By these means, the rapidity and the rate of temperature changes towards reaching the stable value for a given state can be described. The model considers surrounding temperature, CPU

power state, its thermal capacitance and resistance characteristics. One should note that thermal resistance of the circuit may vary in time due to the observed changes in fans' speed. As a result, airflow (distribution and volume) inside a server alters, affecting convective part of thermal resistance, according to:

$$R(t) = R_{cond} + R_{conv}(t) = R_{cond} + \frac{1}{k_v V(t)^n} \quad (3)$$

More details concerning the particular parameters can be found in [11] and [12].

Simultaneously, there is a need to calculate temperature of the air leaving heat source. On one hand, it contributes to the inlet temperature of adjacent components of the server (and thus its ambient temperature). On the other, it constitutes the temperature of (server) outlet air, and thus affects temperature inside server room. This outlet temperature can be calculated in the following way:

$$T_{out}(t) = \frac{Q(t)}{\rho V(t) C_p} + T_{in}(t) \quad (4)$$

By utilising this formula it is possible to describe the mutual impact of neighbouring nodes.

In addition to the thermal models, the power behaviour of particular components will be described. These power models will consider temperature-to-power and energy-to-heat dependencies. They are introduced in [4].

Applying these models will allow characterising thermal state of the server. Moreover, the RECIPE project will benefit from the existing models of data centre, with particular emphasis on cooling facilities [4, 5]. In general, these and similar models, rely on the power characteristics of data centre equipment and the temperature data. Thus, it is essential to capture thermal behaviour of the server in the most precise manner. However, as said above, obtaining these models is an intermediate step in the analyses of the whole system. Moreover, the system will be evaluated against various workload and resource management policies making the thermal state of a single server a subject of dynamic changes. As a result, these models should ensure a proper trade-off between accuracy and performance. They also need to be easily applicable in a large-scale simulation process. One of the possible frameworks supporting such analysis is DCworms, which is described in the next subsection in more detail.

3 Thermal Modelling Tools

3.1 DCworms

DCworms [8] stands for Data Center Workload and Resource Management Simulator and is a simulation toolkit developed by PSNC. It is aimed at simulation of computing infrastructures to estimate their performance, energy consumption and energy-efficiency metrics for different applications and management strategies. DCworms allows modelling the IT resources starting from a single CPU, through servers up to the whole data centre. Additionally, DCworms supports the definition of non-IT devices like fans, CRAHs and chillers. With each resource, it is possible to bind special plugins describing its thermal and power behaviour. These plugins allow estimating the temperature and the power of the given resource with respect to its current state and load profile. Apart from resource modelling, DCworms provides the user flexibility in defining the application structure including various shapes and levels of complexity of workloads combined with time dependencies and preceding constraints between particular jobs and tasks. A dedicated performance plugin allows specifying the behaviour of the particular application with respect to the resources that are assigned to run it. Finally, the user can provide DCworms with the management plugin responsible for executing given scheduling policy and control actions on the defined infrastructure. The outcome of the simulator is a set of statistics and metrics characterising the efficiency and performance of the evaluated environment. More details about the simulator can be found in [8].

4 Summary

The proliferation of heterogeneous HPC systems and applications in new domains has led to new requirements related to non-functional requirements (time, power, reliability, temperature, etc) and the need for platforms to satisfy them. In this deliverable, we provided a summary of the progress achieved in the following activities related to different QoS aspects, as part of WP3:

- Reliability methodology. We reported the methodology used. We showed that it is applicable to any heterogeneous system and we showed how we can integrate all reliability and degradation measurements into a single system value. We described the predictive mechanism we will integrate in the run-time manager.
- Timing analysis. Timing predictions are key to make an efficient use of resources while ensuring that each application is allocated enough resources to complete in time. We reported the mechanism to provide statistically significant timing predictions. This method is already encapsulated in a library and ready for integration in the run-time manager.
- Thermal modelling. Thermal modelling techniques to estimate chip temperature are required in order to enable the assessment of temperature-related reliability effects on chips. The models developed provide the means to know chip temperature accurately with the goal of proactively enhancing chip reliability when no hardware support (i.e. temperature reading) is available.
- Applications characterization. This task has just started (1st of October), so there is nothing relevant to report yet. Details on this topic will be provided in D3.6 in month 30 (October 2020).
- RTRM Reliability, Timing, and Thermal Policies Development. We provide the libraries defined that will serve as an interface with the models developed (task 3.1 to task 3.3). This work is at a preliminary stage, since it spans from April 2019 to March 2021. Yet, the interface has already been defined so we are on track.

This deliverable has described the current status of WP3. Novel contributions on timing, reliability and thermal models have been proposed and submitted for publication. The tasks are progressing according to the plan and we look forward to the use of the novel approaches for deriving better run time manager policies.

References

- [1] G. Agosta, W. Fornaciari, G. Massari, A. Pupykina, F. Reghenzani, and M. Zanella. Managing Heterogeneous Resources in HPC Systems. In *Proc. of PARMA-DITAM '18*, pages 7–12. ACM, 2018.
- [2] A. K. Coskun, D. Atienza, T. Simunic Rosing, T. Brunschweiler, and B. Michel. Energy-efficient variable-flow liquid cooling in 3d stacked architectures. In *2010 Design, Automation Test in Europe Conference Exhibition (DATE 2010)*, pages 111–116, March 2010.
- [3] Ayse Kivilcim Coskun, Tajana Simunic Rosing, and Kenny C. Gross. Temperature management in multiprocessor socs using online learning. In *Proceedings of the 45th Annual Design Automation Conference*, DAC '08, pages 890–893, New York, NY, USA, 2008. ACM.
- [4] Georges Da Costa, Ariel Oleksiak, Wojciech Piatek, Jaume Salom, and Laura Sisó. Minimization of costs and energy consumption in a data center by a workload-based capacity management. In *E2DC*, pages 102–119, 2014.
- [5] Leandro F. Cupertino, Georges Da Costa, Ariel Oleksiak, Wojciech Piatek, Jean-Marc Pierson, Jaume Salom, Laura Sisó, Patricia Stolf, Hongyang Sun, and Thomas Zilio. Energy-efficient, thermal-aware modeling and simulation of data centers: The coolemall approach and evaluation results. *Ad Hoc Networks*, 25:535–553, 2015.
- [6] J. Flich, G. Agosta, et al. Exploring manycore architectures for next-generation HPC systems through the MANGO approach. *Microprocessors and Microsystems*, 61:154 – 170, 2018.
- [7] J. Flich et al. Enabling HPC for QoS-sensitive applications: The MANGO approach. In *2016 Design, Automation Test in Europe Conference Exhibition (DATE)*, pages 702–707, March 2016.
- [8] Krzysztof Kurowski, Ariel Oleksiak, Wojciech Piatek, Tomasz Piontek, W. Przybyszewski, Andrzej, and Jan Weglarz. Dcworms - a tool for simulation of energy efficiency in distributed computing infrastructures. *Simulation Modelling Practice and Theory*, 39:135–151, 2013.
- [9] Clemens J.M. Lasance. Thermally driven reliability issues in microelectronic systems: status-quo and challenges. *Microelectronics Reliability*, 43(12):1969 – 1974, 2003.
- [10] Giuseppe Massari, Anna Pupykina, Giovanni Agosta, and William Fornaciari. Predictive resource management for next-generation high-performance computing heterogeneous platforms. In *Proceedings of the 18th International Conference on Embedded Computer Systems: Architectures, Modeling, and Simulation (SAMOS'19)*, Jul 2019.
- [11] Wojciech Piatek, Ariel Oleksiak, and Georges Da Costa. Energy and thermal models for simulation and workload and resource management in computing systems. *Simulation Modelling Practice and Theory*, 58:40–54, 2015.
- [12] Wojciech Piatek, Ariel Oleksiak, and Micha vor dem Berge. Modelling impact of power-and-thermal-aware fans management on data center energy consumption. *e-Energy*, pages 253–258, 2015.

-
- [13] A. Pupykina and G. Agosta. Optimizing Memory Management in Deeply Heterogeneous HPC Accelerators. In *2017 46th Int'l Conf on Parallel Processing Workshops (ICPPW)*, pages 291–300, Aug 2017.
 - [14] Arvind Sridhar, Alessandro Vincenzi, David Atienza, and Thomas Brunschwiler. 3d-ice: A compact thermal model for early-stage design of liquid-cooled ics. *IEEE Transactions on Computers*, 63(10):2576–2589, 2014.
 - [15] Federico Terraneo, Alberto Leva, and William Fornaciari. An open-hardware platform for mpsoe thermal modeling. In *International Conference on Embedded Computer Systems*. Springer, 2019.
 - [16] Yun Xiang, Thidapat Chantem, Robert P. Dick, X. Sharon Hu, and Li Shang. System-level reliability modeling for mpsoes. In *Proceedings of the Eighth IEEE/ACM/IFIP International Conference on Hardware/Software Codesign and System Synthesis, CODES/ISSS '10*, pages 297–306, New York, NY, USA, 2010. ACM.