Available online at www.prace-ri.eu

Partnership for Advanced Computing in Europe

CFD Simulations in HPC Immersion Cooled

G. Houzeaux[a*], J. Pita[b], O. Lehmkuhl[a], S. Gómez[a], D. Pastrana[a]

[a]*Barcelona Supercomputing Center, C/Jordi Girona 29, 08034-Barcelona, Spain*
[b]*SUBMER Immersion Cooling, Carrer de Pablo Iglesias, 56, 08908-L'Hospitalet de Llobregat, Spain*
guillaume.houzeaux@bsc.es, oriol.lehmkuhl@bsc.es,samuel.gomez@bsc.es, ricard.borrell@bsc.es, dan.ger.pastrana@gmail.com, jaime@submer.com

Abstract

SUBMER company designs, builds and installs Liquid Immersion cooling solutions for HPC, hyperscaler, datacenters, Edge, AI, deep learning and blockchain applications. SUBMER Immersion Cooling is changing how datacenters are being built from the ground up, to be as efficient as possible and to have little or positive impact on the environment around them (reducing their footprint and reducing their consumption of precious resources such as water). SUBMER was founded in 2015 by Daniel Pope and Pol Valls, to make operating and constructing datacenters more sustainable. This SHAPE project aims to increase the knowledge and expertise of SUBMER employees in the realm of CFD simulations in an HPC environment. The project also aims to demonstrate tangible benefits in terms of improved design, efficiency and effectiveness of their immersion cooling solution. The vision of SUBMER is that being able to model a reasonably close estimation to the operational behavior of their systems is key to avoiding costly mistakes, provide a better product to their clientele and greatly reducing time to market. This white paper presents the methodology employed during the project, as well as some selected results.

1. Introduction

SUBMER[†], the clean tech company, designs, builds and installs Liquid Immersion cooling solutions for HPC, hyperscaler, datacenters, Edge, AI, deep learning and blockchain applications. SUBMER Immersion Cooling is changing how datacenters are being built from the ground up, to be as efficient as possible and to have little or positive impact on the environment around them (reducing their footprint and reducing their consumption of precious resources such as water). SUBMER was founded in 2015 by Daniel Pope and Pol Valls, to make operating and constructing datacenters more sustainable.

This SHAPE[‡] project aims to increase the knowledge and expertise of SUBMER employees in the realm of CFD simulations in an HPC environment. The possibility to use a High Performance simulation code and to execute simulations on a supercomputing platform, while being accompanied by HPC experts, is of the utmost importance for the company. The project also aims to demonstrate tangible benefits in terms of improved design, efficiency and effectiveness of their immersion cooling solution. The vision of SUBMER is that being able to model a reasonably close estimation to the operational behavior of their systems is key to avoiding costly mistakes, provide a better product to their clientele and greatly reducing time to market.

This white paper presents part of the work carried out by BSC and SUBMER to achieve the original objectives. Several steps have been undertaken to ensure the minimization of the cost of the simulations while ensuring quality and accuracy of the results. To this aim, the numerical model has been calibrated in terms of: governing equations, geometry simplification, mesh size and boundary conditions. Some modeling has been introduced as well to account for the complex heatsinks of the CPUs and GPUs. After a presentation of the modeling approach, this paper presents the results as well as the introduced modeling assumptions in a chronological way, respecting the timeline of the project.

* Corresponding author. E-mail address: guillaume.houzeaux@bsc.es
† https://submer.com
‡ https://prace-ri.eu/hpc-access/shape-access/

2. Modelling approach

The main computational objective of the project was the simulation of a complete cooling system in operating conditions, including 22 boards of CPUs and GPUs. To achieve this, different intermediate activities have been undertaken. They will be described in the current white paper as they were chronologically carried out:

- First, CPUs and GPUs were characterized individually, and a mesh convergence analysis was conducted, as presented in Sections 3.1 and 3.2.
- Then, we developed a heat source and isotropic porosity models for the heat and the Navier-Stokes equations, respectively, to model the effects of the CPUs, GPUs and heat dissipaters. The goal of such approximation is to greatly simplify the meshing and reduce the number of elements and thus the cost of the simulation. See Section 3.3.
- In order the reduce the deficiencies of the isotropic model, an anisotropic porosity model was next considered. Heat source and porosity were then calibrated against the first individual simulations. This model will be presented in Section 3.4.
- We then considered a new board in Section 4 that was validated as well.
- This new board was then placed into the cooling system and duplicated to obtain a complete configuration of the cooling system in Section 5. We then carried out simulations with a 70M mesh, and to ensure mesh convergence, the same simulation on a multiplied mesh (70*8) of 560M elements was conducted. Results have enabled SUBMER to clearly identify the zone where heat is not properly evacuated.

The governing equations are the incompressible Navier-Stokes equations, under the Boussinesq approximation, to account for relatively small temperature changes, justified in the current context:

$$\rho\frac{\partial \boldsymbol{u}}{\partial t} + \rho(\boldsymbol{u}\cdot\nabla)\boldsymbol{u} - \nabla\cdot[2\mu\boldsymbol{\varepsilon}(\boldsymbol{u})] + \nabla p = \rho\beta(T - T_0)$$

$$\nabla\cdot\boldsymbol{u} = 0$$

$$\rho C_p\frac{\partial T}{\partial t} + \rho C_p\boldsymbol{u}\cdot\nabla T - \nabla\cdot(k\nabla T) = 0$$

The physical properties are given in the following table:

Property	Symbol	Value	Units
Density	ρ	780	kg/m^3
Viscosity	μ	468 x 10^{-5}	kg/(ms)
Thermal conductivity	k	0.14	W/(mK)
Specific heat	C_p	2200	J/(kg K)
Thermal expansion coefficient	β	0.0007	K^{-1}
Reference temperature	T_0	308.15	K

The code used in this project to carry out the CFD simulations was Alya[§]. Alya is developed at BSC and is part of the Unified European Applications Benchmark Suite (UEABS[**]) of PRACE. The fluid equations are solved with a Finite Element variational multiscale method, with in addition convection and time tracking of the subgrid scale, as detailed in [HP08] These equations are solved with an implicit method presented in [HA11] . The iterative solver for the pressure equation is presented in [LM11]. The solution of the temperature equation is implicit as

§ https://gitlab.com/bsc-alya/alya

** https://repository.prace-ri.eu/git/UEABS/ueabs/

well, and an algebraic subgrid scale method is considered for stabilization purposes [HB09]. These equations are solved in parallel, using MPI for communication. The parallelization is extensively described in [VH16].

We will start by carrying out simulations on two configurations of a first board, referred to as board 1. On the one hand, a single CPU series of numerical tests will show mesh convergence and study different approximations. On the other hand, we will consider the complete board simulation. The meshes will be referred to MMx, where x is the level of mesh multiplication, the original mesh being referred to as MM0.

3. Board 1

3.1. CPU characterization

Several approximations have been considered for this first set of simulations: pure conduction, incompressible Navier-Stokes equations (NSI) and Boussinesq approximation (NSB) as given by the previous set of equations. Details of the original mesh, referred to as MM0 are shown in the Figure 1. In order to perform the mesh convergence, we have used the mesh multiplication described in [HC13] .

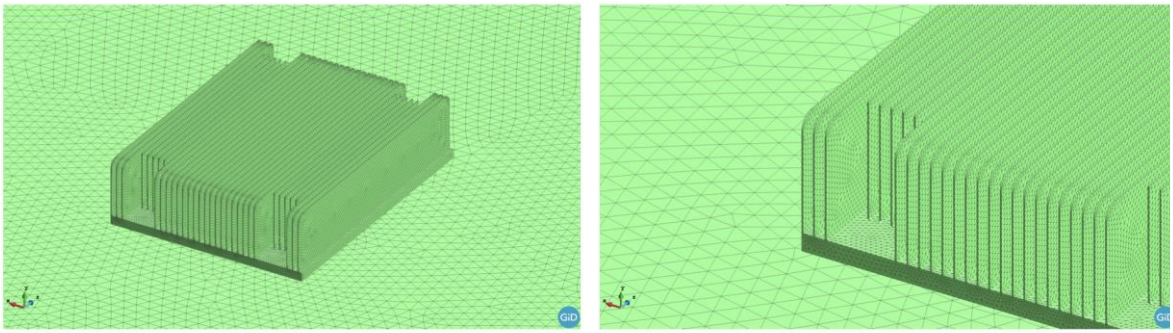

Figure 1 : Details of the MM0 mesh used for CPU characterization.

Two meshes have been generated, MM0 and MM1, with 4.1M elements and 32.7M elements, respectively. Pure conduction results are not shown here as these simulations were only conducted to have an insight on mesh convergence. We will first compare the evolution of the maximum temperature using NSI and NSB, shown in Figure 2.

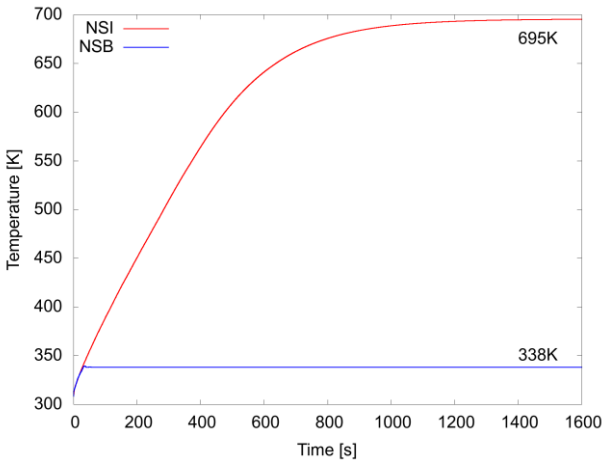

Figure 2 : Evolution of maximum temperature for NSI and NSB

We observe a huge difference in maximum temperature. By looking at the flow inside the dissipator, Figure 3, we observe that in the case of NSI, the flow is slowed down very quickly when entering the dissipator. If z is the direction of the flow, this means that $du/dz < 0$. In order to satisfy the incompressibility conditions, the flow has to rise up, that is $dv/dy > 0$, where y is the vertical direction. This upward flow prevents the heat generated right after the upward motion to be dissipated. In the case of NSB, natural convection aligns the flow along the dissipator channels, which enables to evacuate the heat much more efficiently. These effects are reflected in the difference in maximum temperature noted in Figure 2.

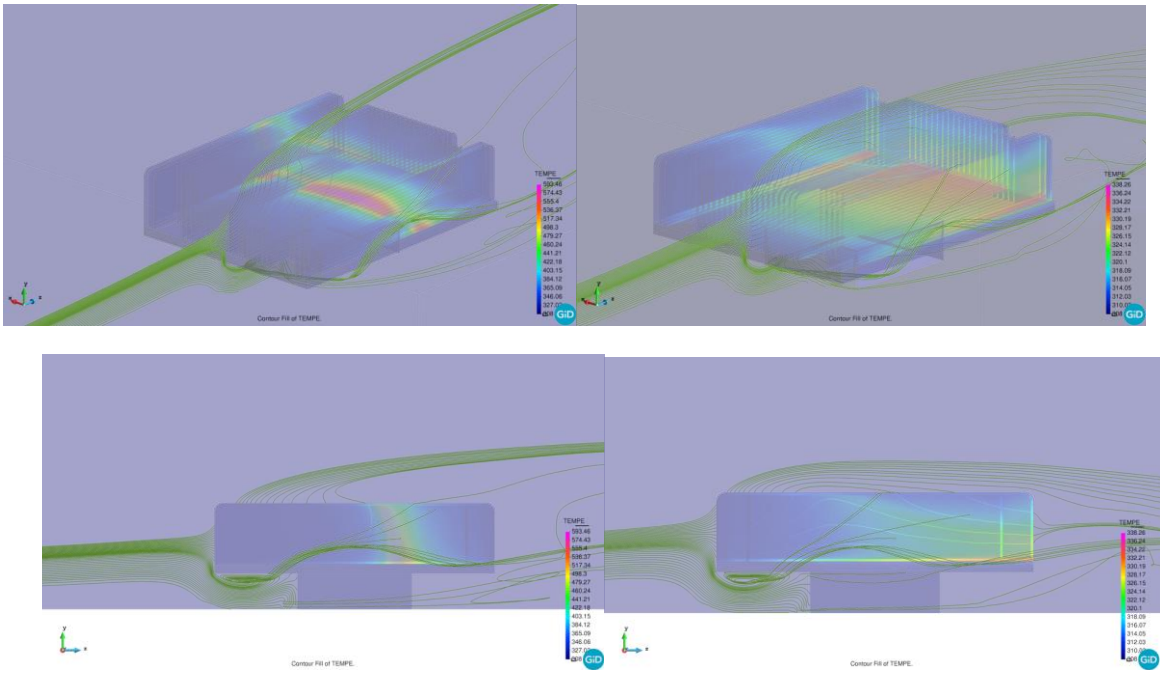

Figure 3: Streamlines and temperature contours. Hot spots (red and magenta) are over 330K and cool spots (blue) are around 310K. (left) NSI, flow is rapidly slowed down and moves upward. (right) NSB, flow is much more aligned with the dissipator channels due to natural convection.

Finally, Figure 4 shows the evolution of maximum temperature and average temperature on the dissipator, for the two meshes MM0 and MM1. Small differences are observed for these global values.

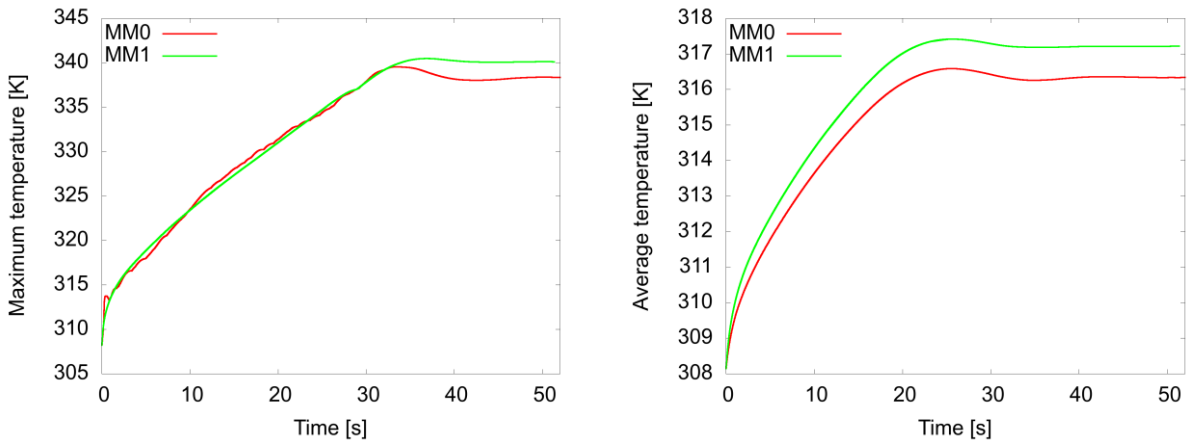

Figure 4 : NSB, mesh convergence. (Left) Maximum and (Right) average temperatures on dissipator.

3.2. Complete board

The specifications of CPUs and GPUs are given in the next Table.

Device	Model	Power [W] per unit	Number of units
CPU	Intel Xeon E5-2697 v4	145	2
GPU	Nvidia Tesla K80	300	8

The final mesh is composed of 34.6M tetrahedra elements, CPU and GPU elements have the resolution of the MM0 isolated CPU characterization experiment (previous section). Special refinements zones are located also at the processor zones to substitute the small grating elements of the real geometry. Some details of the mesh are depicted in Figure 5, where a general view of the mesh at the center plane is presented together with a detail of the internal mesh of the CPU dissipator. The mesh has been constructed to have good aspect ratio with growth ratio of 1.2.

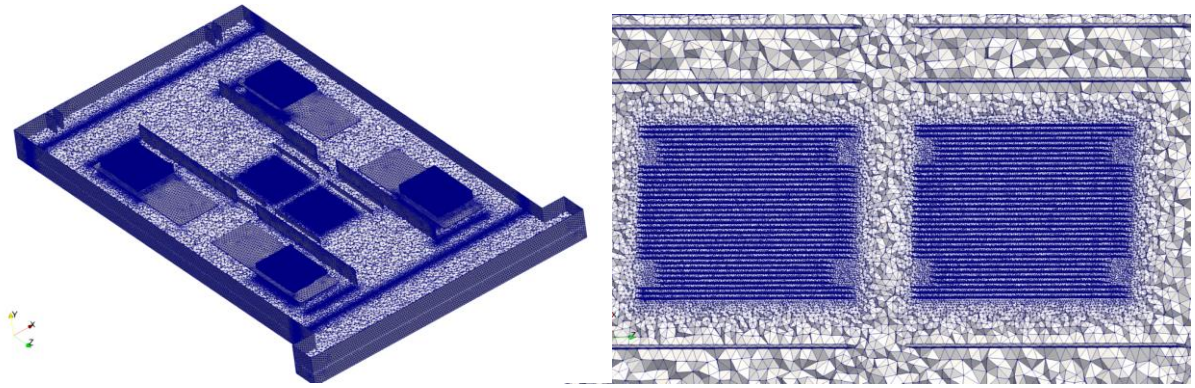

Figure 5: Slice of the 3D volumetric mesh at the center plane and mesh detail at the CPUs dissipator.

In order to have a first glance over the general flow topology the reader is directed to Figure 6 left, where 2D streamlines at the center plane of the server are presented.

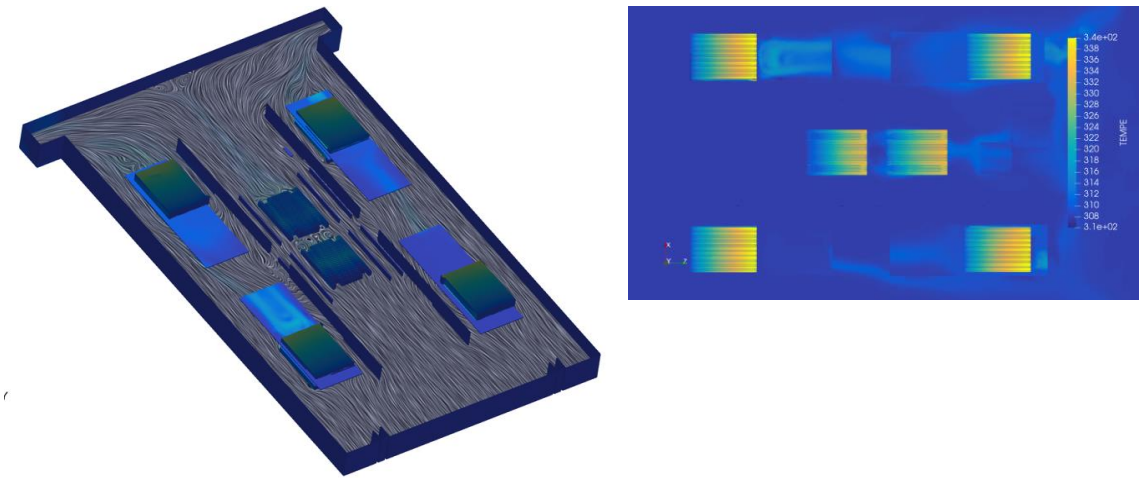

Figure 6 : Global flow 2D streamlines and skin temperature on devices.

It's worth to notice that the flow remains well aligned with the main flow direction, mainly by the strong effect of the buoyancy force. The only disturbances of the flow are due to the recirculation bubbles generated by the CPUs and GPUs. Of special interest is the CPU zone, where a strong recirculation after the first one heavily affects the thermal behaviour of the next CPU. This is clearly shown in Figure 6 right, where skin temperature for all the relevant surfaces are presented. Here one can see how the first CPU behaves very similar to an isolated CPU, however in the case of the second CPU is hotter from the beginning of the heat exchanger. The reason is the strong effect of the flow recirculations on the incident liquid, thus, blocking the entrainment of fresh liquid into the CPU dissipator (see Figure 8). Finally, flow characteristics over the GPU elements can be assessed in Figure 7 where 3D streamlines and temperature iso-surfaces are presented. Although the flow file looks symmetrical, an asymmetry in the thermal field is observed. However, the thermal behaviour is good and all the components are correctly cooled.

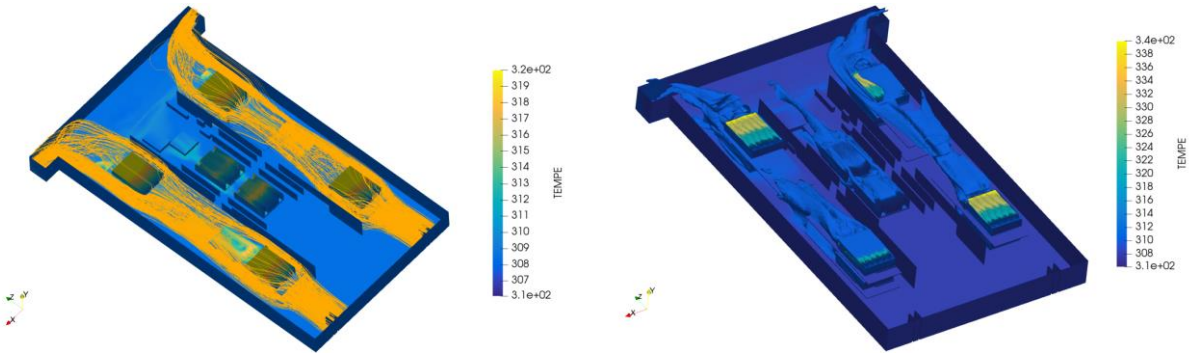

Figure 7: Skin temperatures with 3D streamlines over the GPUs and temperature iso-surfaces (lv: 310-320-330-340-350).

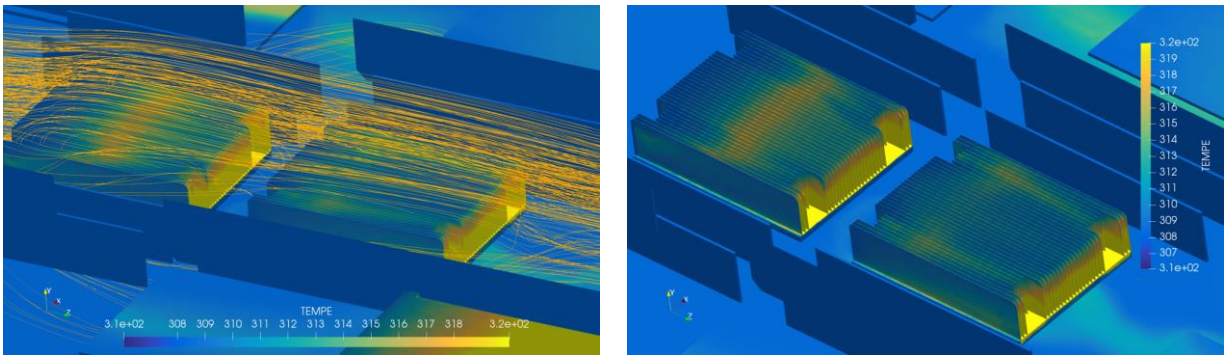

Figure 8: 3D streamlines over the CPUs region and CPUs dissipator skin temperature.

3.3. Reduced isotropic material model

In order to reduce the computational cost of the simulations, a reduced model for the heatsinks physical characteristics is proposed. This model consists in the replacement of the detailed heat sinks geometry by materials defined as regions on the computational domains with a defined isotropic porosity σ and a volumetric heat source parameter *Q*. We will consider a single heat sink for the GPU and two different heat sinks (heat sink1 and heat sink2) for the CPUs to account for the upstream and downstream CPUs, respectively. The governing equations thus become:

$$\rho\frac{\partial \boldsymbol{u}}{\partial t} + \rho(\boldsymbol{u}\cdot\nabla)\boldsymbol{u} - \nabla\cdot[2\mu\boldsymbol{\varepsilon}(\boldsymbol{u})] + \sigma\boldsymbol{u} + \nabla p = \rho\beta(T - T_0)$$

$$\nabla\cdot\boldsymbol{u} = 0$$

$$\rho C_p\frac{\partial T}{\partial t} + \rho C_p\boldsymbol{u}\cdot\nabla T - \nabla\cdot(k\nabla T) = Q$$

The porosity was estimated from the complete board simulations results described in the previous section. Average values for the pressure drop and the inlet velocity on the heat sinks were computed to determine the porosity. Figure 9 shows the values used to calibrate the GPU porosity. Similar results were also obtained for the CPU.

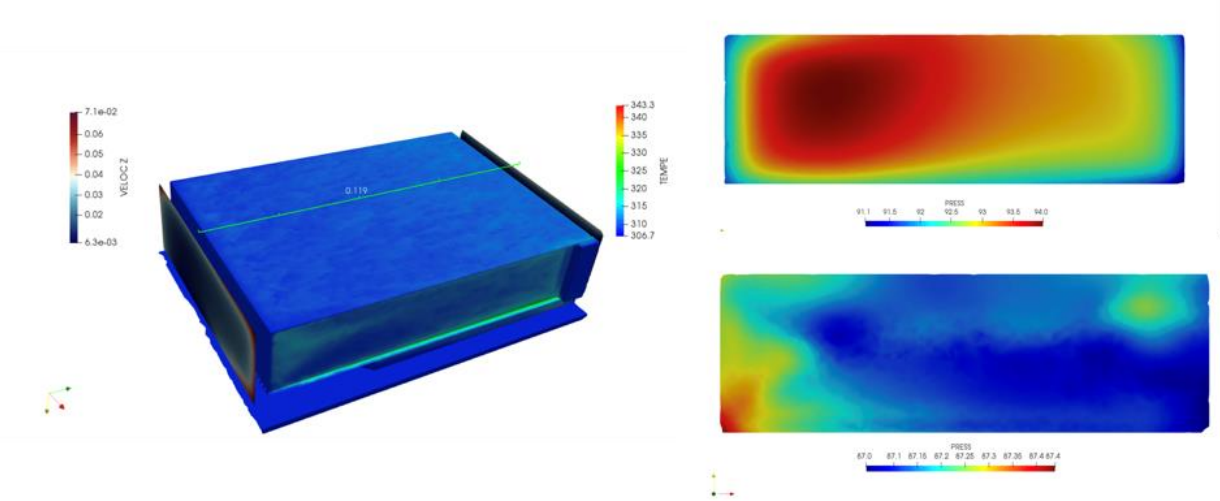

Figure 9 : (left) Velocity at the inlet and skin temperature; (right) pressure at front and back of the GPU heat sink. These results were obtained to calibrate the porosity model.

After calibrating the model, the reduced models was then compared to the ones obtained on the detailed geometries. Figure 10 and Figure 11 show a comparison between the reduced material model and the detailed geometry for the CPU heat sinks1 and CPU heat sinks2, respectively. Here, the heatsink midplane temperature distribution and the average value of the temperature over the heatsink fluid volume T_m are shown. Similar results have been obtained for the GPU as well.

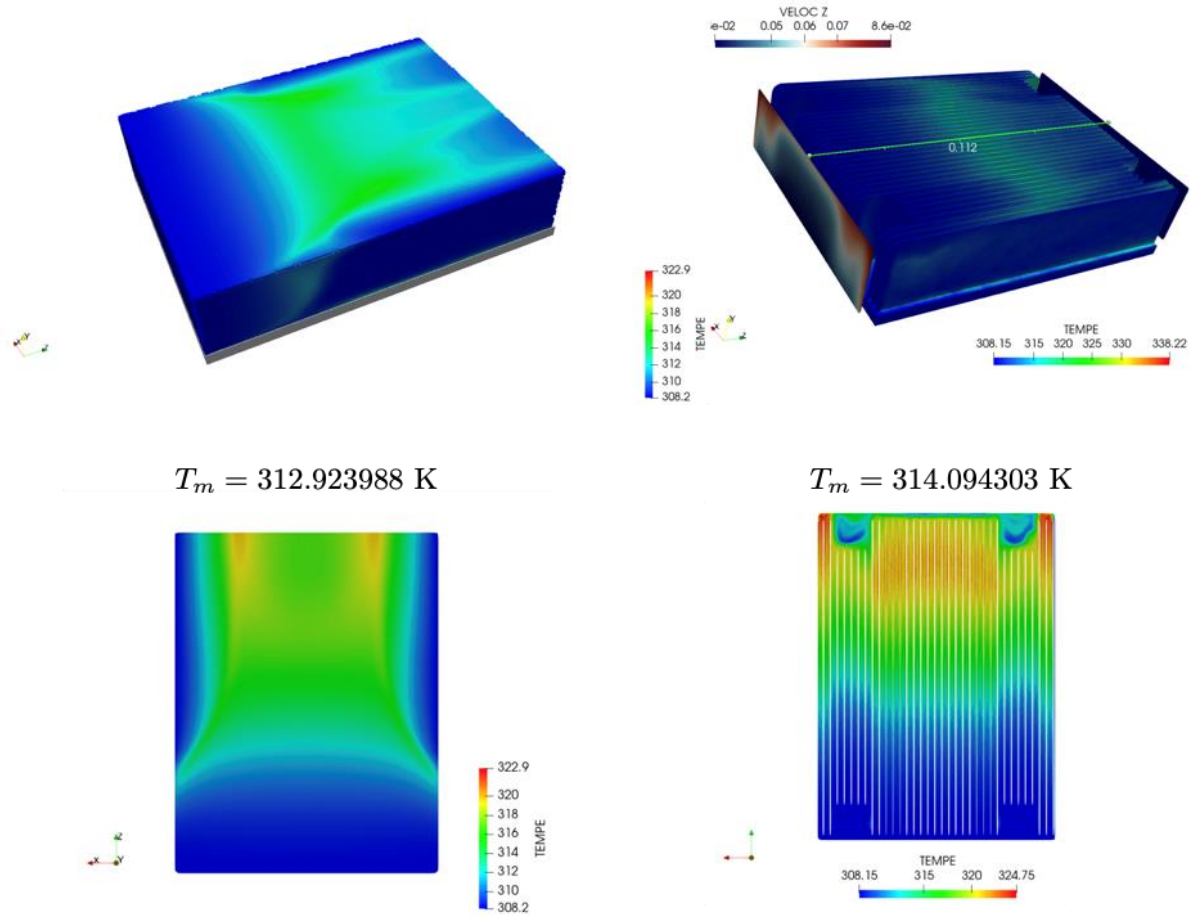

Figure 10: Comparison between the temperature distribution for CPU heat sink1: (left) reduced material model with isotropic porosity; (right) detailed heatsink.

We note that the resemblance of the temperature distribution is remarkable. It has to be taken into account that the porosity is isotropic, this is why there is a component of the velocity in the x-direction at the sides of the material model that drives the temperature towards the center of the heatsink. This effect is more important for the CPU-heat-sink1 (Figure 10). Also, it is important to note that the difference in the maximum of skin temperature and the average temperature over the fluid volume T_m between the material reduced model and the detail geometry heatsink is higher for the CPU heat sink1. This seems to point out that the CPU heat sink2 also has an effect on CPU heat sink1 which cannot be reproduced in the material reduced model isolated simulation. For the CPU heat sink2 (Figure 11) the average temperature over the heatsink fluid volume T_m is quite similar. In this case, the effect of the CPU heat sink1 is taken into account just by the reduction of the inlet velocity.

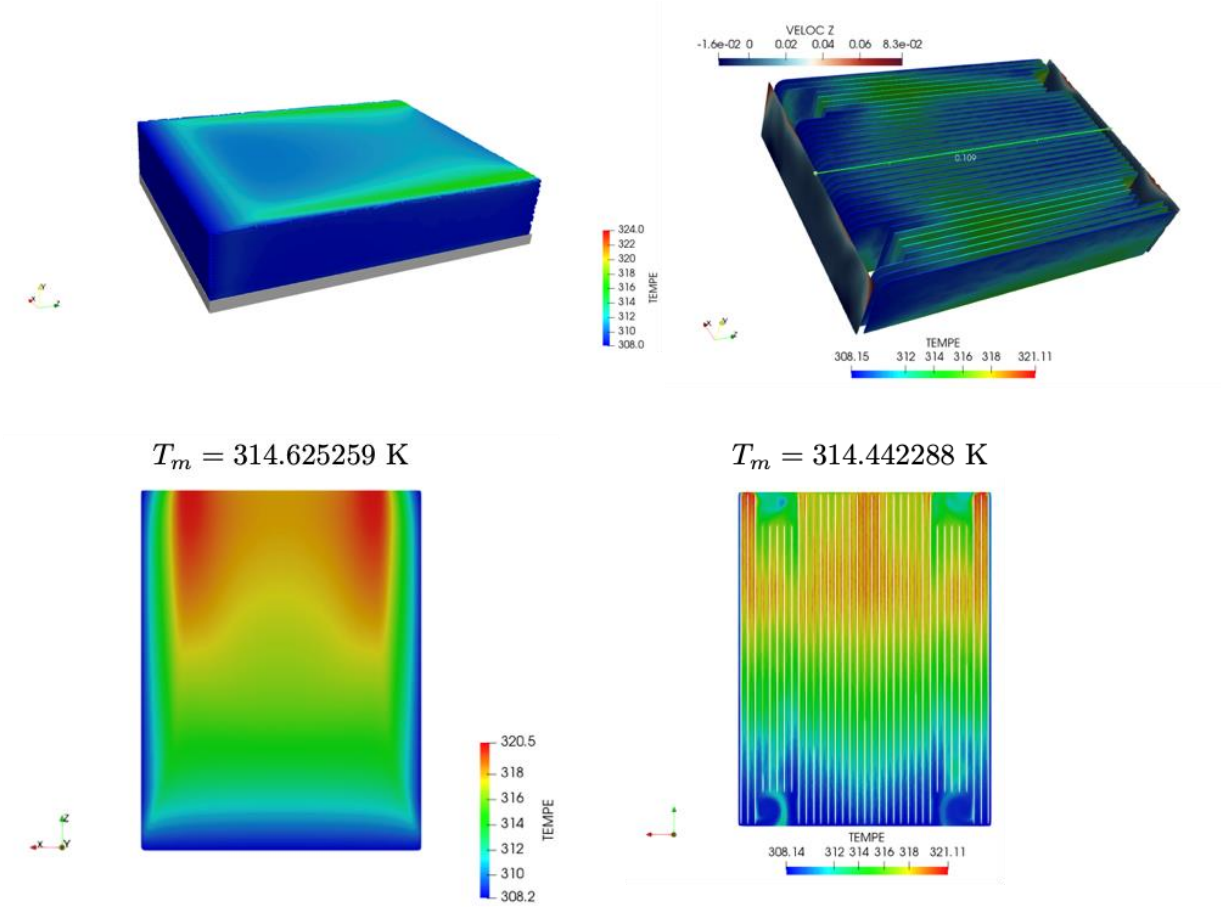

Figure 11: Comparison between the temperature distribution for CPU heat sink2. (left) reduced material model with anisotropic porosity; (right) detailed heat sink.

3.4. Reduced anisotropic material model

A more accurate approach for the reduced material model consists in using an anisotropic porosity in order to prescribe a high porosity value in the crossflow direction x and a lower one in the other two directions. Now the porosity $\boldsymbol{\sigma}$ is a 3x3 tensor which carries the directional information of the porosity as:

$$\rho\frac{\partial \boldsymbol{u}}{\partial t} + \rho(\boldsymbol{u}\cdot\nabla)\boldsymbol{u} - \nabla\cdot[2\mu\boldsymbol{\varepsilon}(\boldsymbol{u})] + \boldsymbol{\sigma}\boldsymbol{u} + \nabla p = \rho\beta(T - T_0)$$

In Figure 12 the effect of the anisotropic porosity for the CPUs is shown after calibration. The x-direction component of the velocity is mitigated by the high value of the porosity in this direction and an increment in the average temperature over the fluid volume enclosed by the material is obtained for the CPU heat sink1 which now is closer to the temperature of the detailed geometry heatsink (see Figure 10). In the case of the CPU heat sink2, no important differences can be observed but a small reduction of about 0.3 K in the average over the fluid volume is registered.

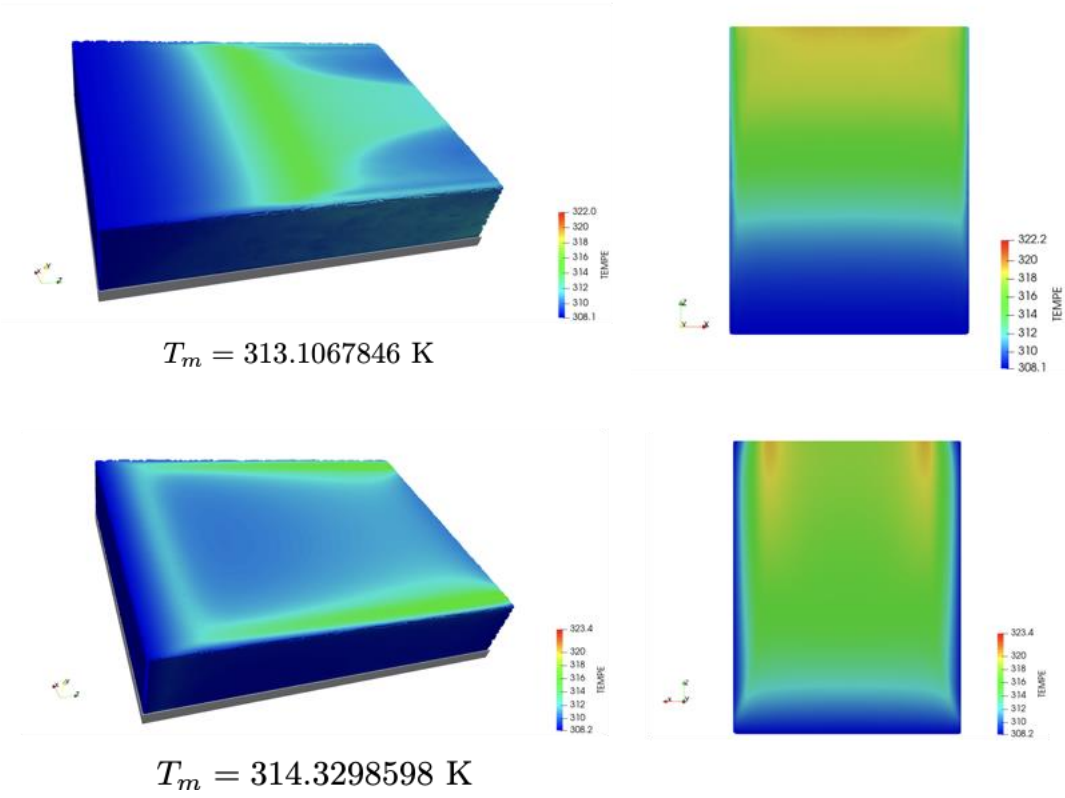

Figure 12: Skin temperature and temperature over a midplane for the reduced material model for CPU heat sink1 (top) and CPU heat sink2 (bottom). T_m is the average temperature over the fluid volume enclosed by the material.

Finally, another simulation with the material including anisotropic porosity and volumetric heat flux parameters computed for the GPU heatsink was performed. For this case, as the GPU heat sinks are surrounded by walls a high porosity value is prescribed not only in the x-direction but also in y. Figure 13 shows a comparison of the temperature distribution over the midplane perpendicular to y. As can be observed, the temperature in the reduced material model is about 2 K higher than the detail geometry GPU heat sink. The same happens with the average temperature of the fluid enclosed by the heatsink and the material T_m.

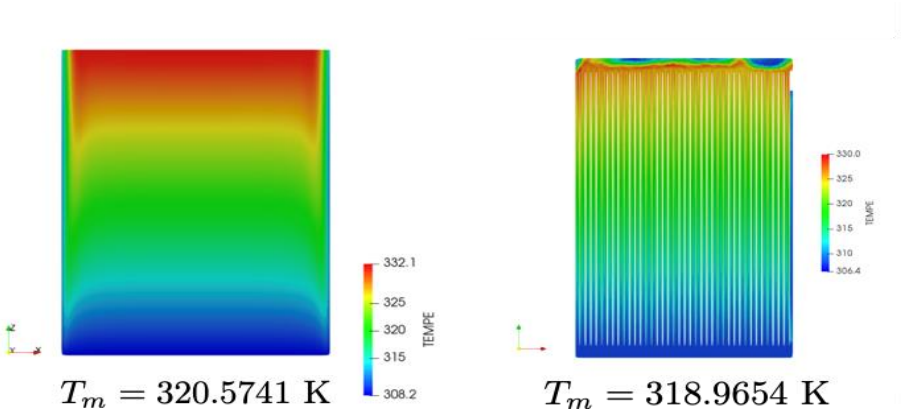

Figure 13: Temperature distribution over the midplane perpendicular to y. (left) reduced material model with anisotropic porosity; (right) detailed heat sink. T_m is the average temperature of the fluid enclosed by the heat sink and the material.

4. Board 2

The second board is the final board to be considered in next section for the complete immersion cooling system. As the components are similar to that of the previous board, it was simulated using the same heat sink and anisotropic porosity. Two meshes have been considered, MM0 and MM1, with 1.90M elements and 15.2M respectively. Figure 14 shows the evolution of the maximum velocity and temperature for the two meshes. We observe very similar behaviors for both quantities.

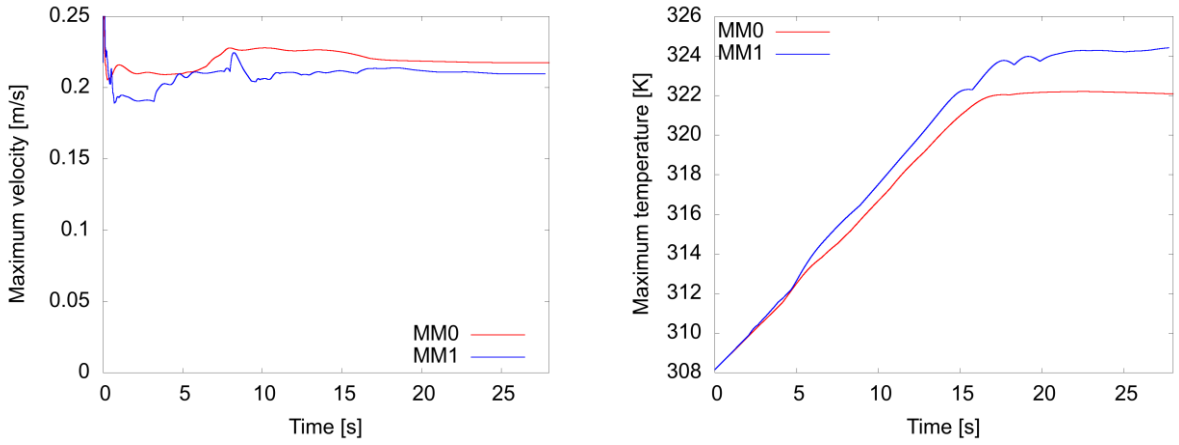

Figure 14: Board 2, mesh convergence: Maximum velocity and temperature.

Results of velocity and temperature distributions are shown in Figure 15.

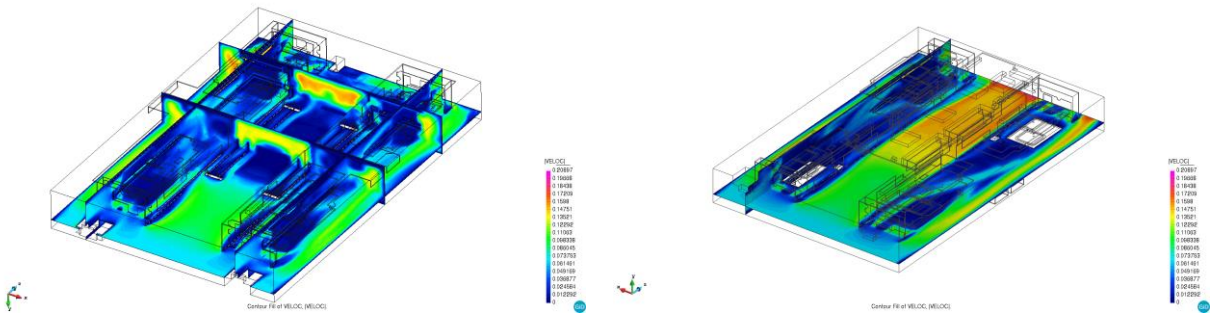

Figure 15: Board 2, MM1: velocity and temperature contours.

5. Immersion cooling system with board 2

A first set of simulations has been carried on a 70M mesh (MM0) to adjust the local mesh size and make sure outflow conditions were imposed sufficiently far to avoid back flow. Then, the simulation of the complete setup using a multiplied mesh MM1 with 560M elements was carried out using the anisotropic porosity model, together with a heat source to model the CPUs and GPUs. The simulation was carried out using 12000 cores for 16h, consisting of 192.000 core hours.

7 degrees difference was observed in the maximum temperature between the MM0 and MM1 meshes, as shown in Figure 16.

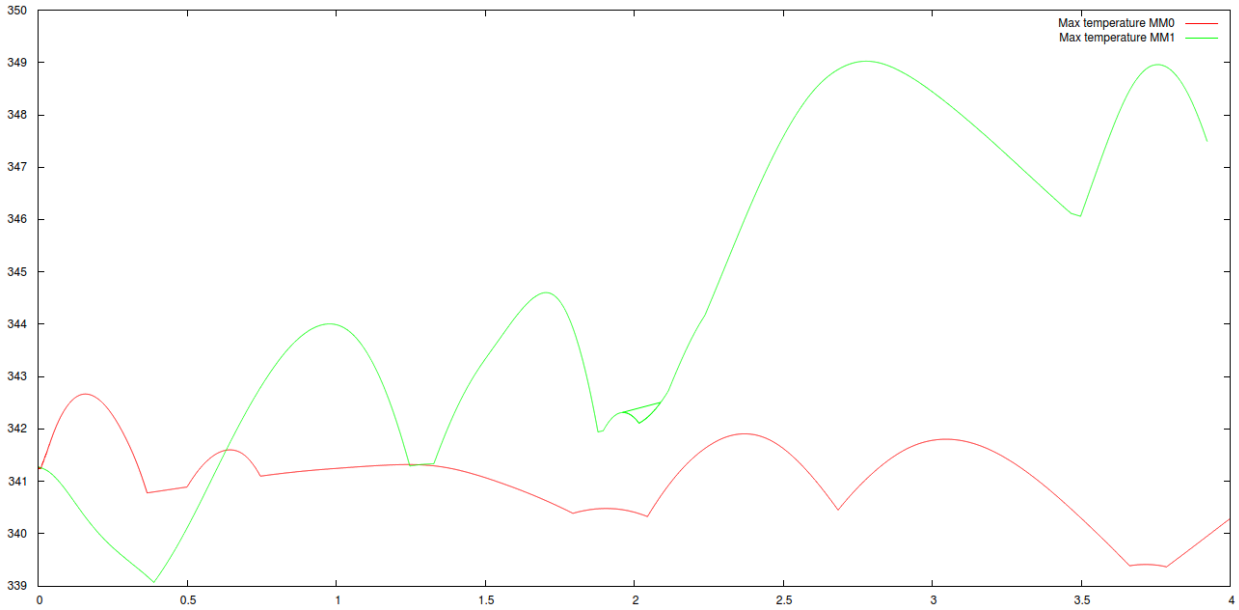

Figure 16: Comparison of maximum temperature for M0 and M1 meshes

For a better understanding of the flow some screenshots of zones of interest are shown. Figure 17, Figure 18 and Figure 19 show the complexity of the flow and temperature profiles inside the setup.

NSB - MM1

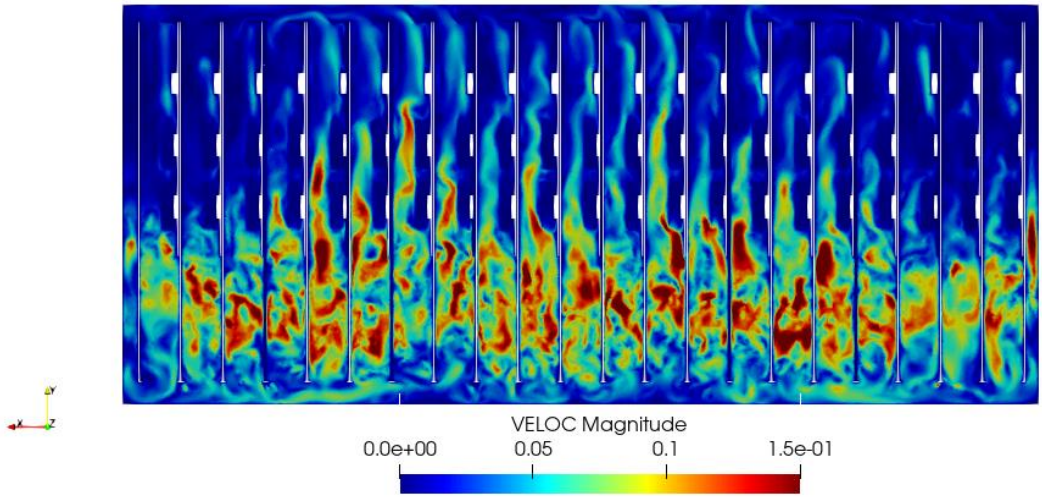

Figure 17. velocity snapshot in the complete setup.

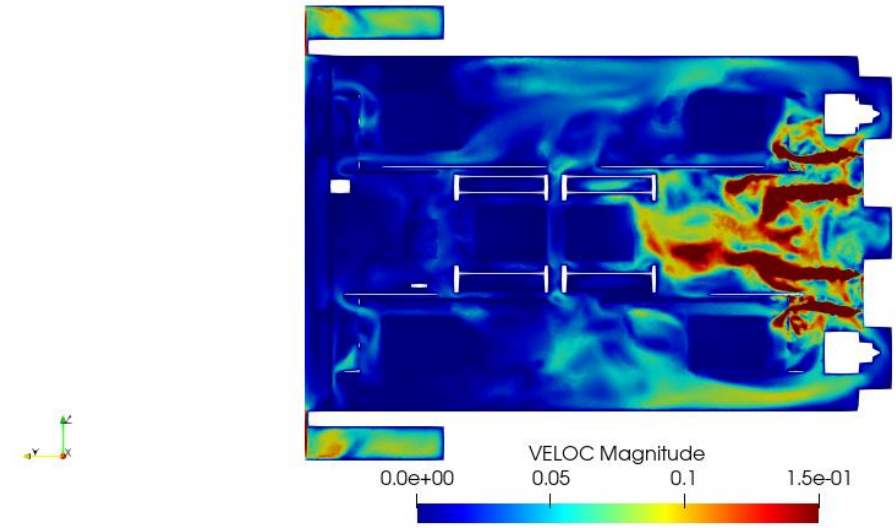

Figure 18: Horizontal plane of s single board.

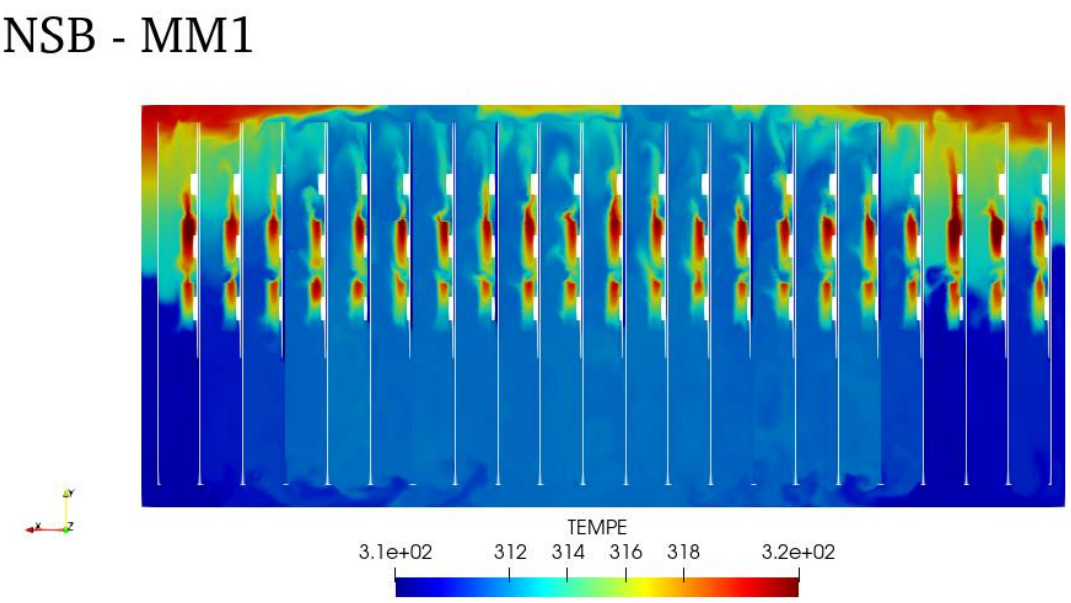

Figure 19: Temperature snapshot in the complete setup.

6. Conclusions and Outlook

This paper is the result of a SHAPE project where our aim was to demonstrate the ability of HPC simulations to predict not only global behaviour but also local details of the flow field. This last capability is important to detect anomalies and to provide insights to the manufacturer to design solutions. In particular, zones of difficult heat evacuation have been identified, not only on single board but also on the complete setup.

Acknowledgements

This work was financially supported by the PRACE-6IP project funded in part by the EU's Horizon 2020 research and innovation programme (2014-2020) under grant agreement INFRAEDI-823767.

References

[HP08] G. Houzeaux and J. Príncipe. A Variational Subgrid Scale Model for Transient Incompressible Flows. Int. J. CFD, 22(3):135-152, 2008.

[HB09] G. Houzeaux, B. Eguzkitza, and M. Vázquez. A Variational Multiscale Model for the Advection-Diffusion-Reaction Equation. CNME, 25(7):787-809, 2009

[HC13] G. Houzeaux, R. de la Cruz, H. Owen, and M. Vázquez. Parallel uniform mesh multiplication applied to a Navier-Stokes solver. Computers & Fluids, 80:142--151, 7 2013.

[HA11] G. Houzeaux, R. Aubry, and M. Vázquez. Extension of fractional step techniques for incompressible flows: The preconditioned Orthomin(1) for the pressure Schur complement. Computers & Fluids, 44:297-313, 2011.

[LM11] Löhner, F. Mut, J. Cebral, R. Aubry, and G. Houzeaux. Deflated Preconditioned Conjugate Gradient Solvers for the Pressure-Poisson Equation: Extensions and Improvements. Int. J. Numer. Meth. Engn., 87:2-14, 2011.

[VH16] M. Vázquez, G. Houzeaux, S. Koric, A. Artigues, J. Aguado-Sierra, R. Arìs, D. Mira, H. Calmet, F. Cucchietti, H. Owen, A. Taha, E. Dering Burness, J. M. Cela, and M. Valero. Alya: Multiphysics engineering simulation towards exascale. J. Comput. Sci., 14:15-27, 2016.