Available online at www.prace-ri.eu

Partnership for Advanced Computing in Europe

Tooling4G – Minimize the Airflow Generated Noise on Automotive HVAC Systems

J. F. Caseiro*[a], P. Alberto[b]

[a]*Technological Center for Mouldmaking Special Tooling and Plastic Industries (CENTIMFE), Marinha-Grande, Portugal*
[b]*Laboratory for Advanced Computing – University of Coimbra, Coimbra, Portugal*

Abstract

The main goal of this work is to assess the gain that can be obtained when considering High-Performance Computing (HPC) for solving complex industrial problems. To that end, a Heat, Ventilation and Air Conditioning (HVAC) system of an automobile is analysed by means of an Aero-Acoustic Computational Fluid Dynamics (CFD) analysis. The performances of the HPC Navigator from the University of Coimbra and a research workstation are compared using Ansys® Fluent. Results show that using the Navigator with 64 cores distributed equally on two nodes reduces by a factor of 12.3 the execution time when compared with a workstation limited to 4 cores. The use of HPC resources is extremely useful for Small and Medium Enterprises (SMEs) that don't have the required hardware to solve this kind of problems, enabling an increase in productivity and reducing costs associated with shorter project stages.

Introduction

With the evolution of vehicular motorization, and in particular the use of electric motors, the noise related to the working engine has been significantly reduced. Consequently, noise sources that might have been considered as minor, are now more relevant. In particular, the Heat, Ventilation and Air Conditioning (HVAC) system of an automobile has a significant impact in the perception of comfort of the users. A positive perception regarding sound quality can provide a competitive advantage for the manufacturer [1].

The present project deals with the numerical simulation, using Computational Fluid Dynamics (CFD), of the airflow in an automotive HVAC system. The assessment of the noise generated by the airflow may be computed by means of Computational Aero-Acoustic models. However, these models require the use of Scale Resolving strategies such as the Large Eddy Simulation (LES) or Detached-Eddy Simulation (DES) which must be combined with very fine meshes, especially near the physical boundaries of the numerical model. In addition, the flow time step required to capture the sound-related frequencies of interest is also small.

[a] Corresponding author email address: joao.caseiro@centimfe.com

All the above-mentioned factors give rise to complex numerical models that may require several weeks of execution time. Therefore, it is intended to demonstrate the advantages of using High-Performance Computing (HPC) to solve such problems using the Navigator Cluster from the University of Coimbra. This study is especially relevant for Small and Medium Enterprises (SMEs) that need, but are not able to have, the required computational power (hardware) in-house.

In the present work, an aero-acoustic analysis of an automotive airvent is tackled. Firstly, the performance of a standard workstation is compared with the Navigator Cluster. The speedup and efficiency of the latter is also assessed for different mesh densities and parameter configurations. Finally, the aero-acoustic simulation is performed, and the obtained results are compared with the available experimental data.

Model Description

The geometry to be analysed corresponds to a simplified automotive airvent, as depicted in Figure 1, comprised of a housing and two sets of lamellae that are used to direct the airflow. The numerical model will then correspond to the interior of the airvent, plus an extension on the outlet that represents a portion of the passenger space (Figure 2). To reduce the problem size, symmetry conditions were considered. The air inflow was taken as 400m^3/h. Further details about the mesh and mathematical models employed are given in the subsequent sections. The numerical simulations were performed using the software Ansys® Fluent [2].

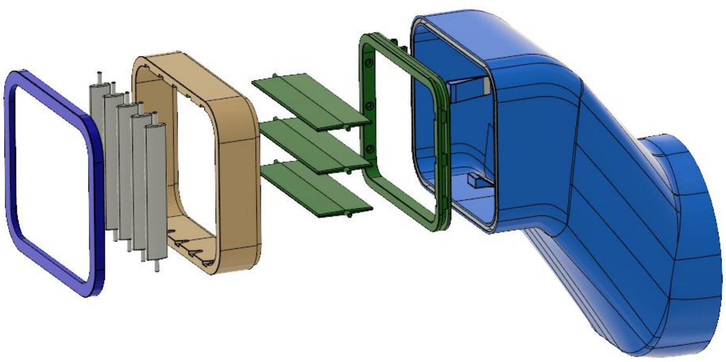

Figure 1 - Geometry of the airvent.

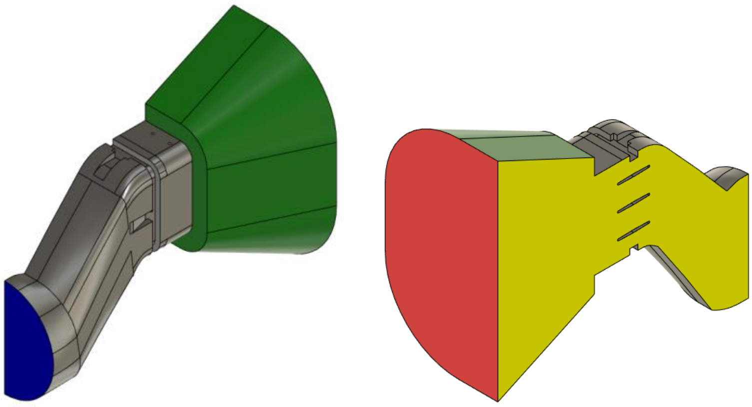

Figure 2 - Interior of the airvent with inlet (blue), outlet (red) and symmetry (yellow) conditions.

The various components of an airvent for the automotive industry are obtained by the Plastic Injection Moulding (PIM) process. A physical prototype of the airvent was printed using Selective Laser Sintering (SLS), polished and painted to obtain a surface finishing equivalent to a component obtained by PIM. Afterwards, the prototype was tested in an anechoic chamber and the data is used to validate the numerical results obtained.

Navigator and Workstation Performance

The first step is to compare the performance of the research workstation available at Centimfe and the HPC Navigator Cluster from the University of Coimbra.
The workstation at Centimfe is composed by an Intel Xeon Gold 6128 CPU @ 3.4GHz with 64GB of RAM. Due to restrictions on the license of the Ansys® Fluent software, the maximum number of cores available for the computations is 4. For the Navigator HPC, the cores used for the simulation belong to a single node on the cpu2 partition with 2 x Intel Xeon Gold 6148 (20-core) @ 2.40 GHz and 96GB of memory. The license on the simulation software only allows for a maximum of 64 cores.

Regarding the numerical model, the turbulence is described by the Realizable k-epsilon model with Scalable Wall Functions. The pressure-velocity coupling uses the SIMPLE scheme. The mesh is composed by approximately 1.7 million cells. Since a 3D model with turbulence is considered, a total of 6 equations (four for pressure and velocity along x, y, and z and two due to turbulence parameters) are solved at each cell at each iteration. For illustration purposes, the mesh is depicted in Figure 3.

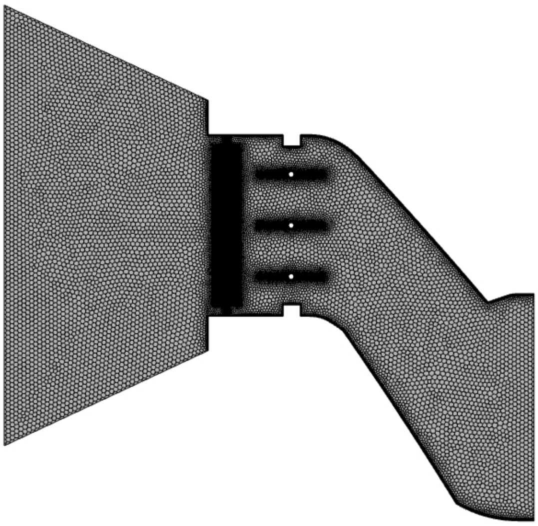

Figure 3 - Airvent interior discretized by 1.7 million cells.

The simulation is divided into 3 steps that correspond to reading the case file containing all the information related to the model (mesh, boundary condition, mathematical models, etc.), followed by an iterative process until convergence is obtained and, finally, the results are written to an output file. The files necessary to perform the simulation were pre-generated in the workstation and uploaded to the Navigator Cluster.

In Figure 4, the computation time for the workstation and the Navigator Cluster is presented. To facilitate the comparison, the results are normalized with respect to the execution time obtained by the workstation using a single core. As can be seen, by using the Navigator, there is a reduction of execution time of about 25, 74 and 87% when considering1, 2 and 4 cores, respectively, when compared with the research workstation with a single core.

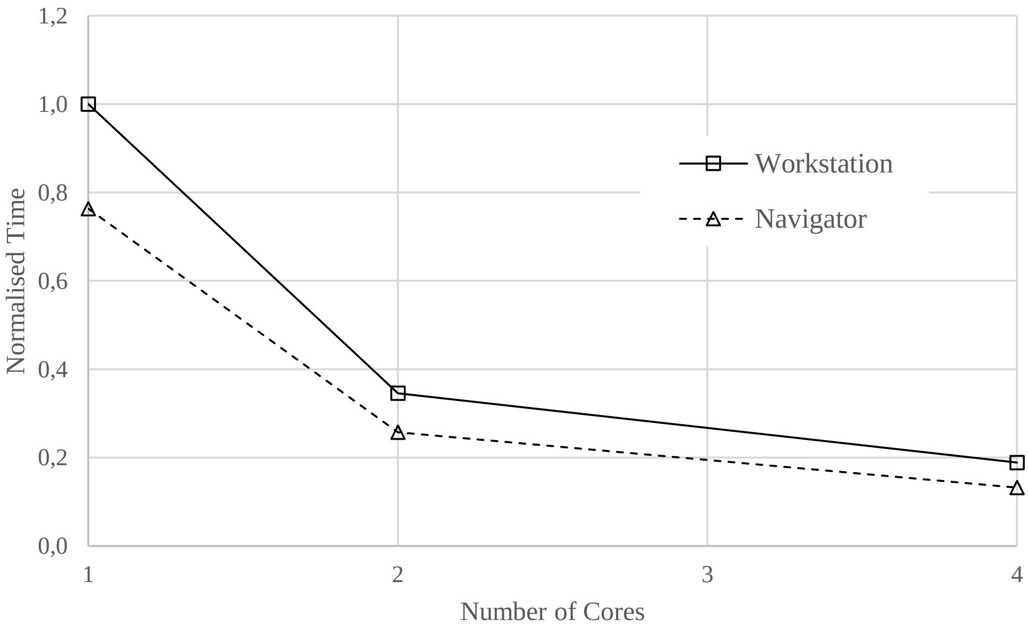

Figure 4 - Normalized execution time for the workstation and Navigator Cluster.

The performance of the Navigator cluster on a single node is further investigated by performing speedup and efficiency analyses. The speedup S(p) is defined as

$$S(p) = \frac{t(1)}{t(p)},$$

where t(1) and t(p) are the execution times with one and p processing units, respectively. The efficiency E(p) is measured as

$$E(p) = \frac{S(p)}{p} = \frac{t(1)}{p \times t(p)}.$$

The speedup and efficiency results can be seen in Figure 5 and Figure 6, respectively. There is a superlinear speedup when moving from 1 to 2 cores, but afterwards the efficiency of the process steadily decreases. Although we did not do a detailed analysis of this behaviour, we believe this was due to the fact that, due to domain decomposition, the memory for each process was reduced and it was possible to use only cache memory, which is faster. This kind of behaviour has been documented before, at least in OpenMP applications. The speedup increases almost linearly up to 32 cores and with 64 it decreases slightly. This last fact is almost certainly due to the use of hyperthreading for 64 processes, since each node has only 40 physical cores able to do floating point operations. When compared to the workstation running of 4 cores, the Navigator Cluster running at 32 cores is 7 times faster.

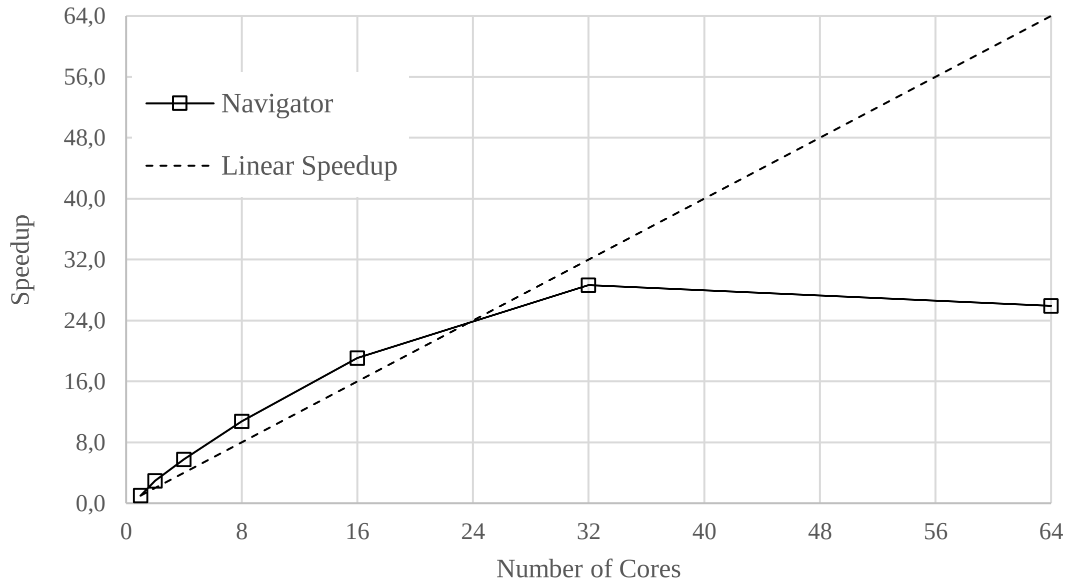

Figure 5 - Speedup for a 1.7 million cell mesh considering a single node.

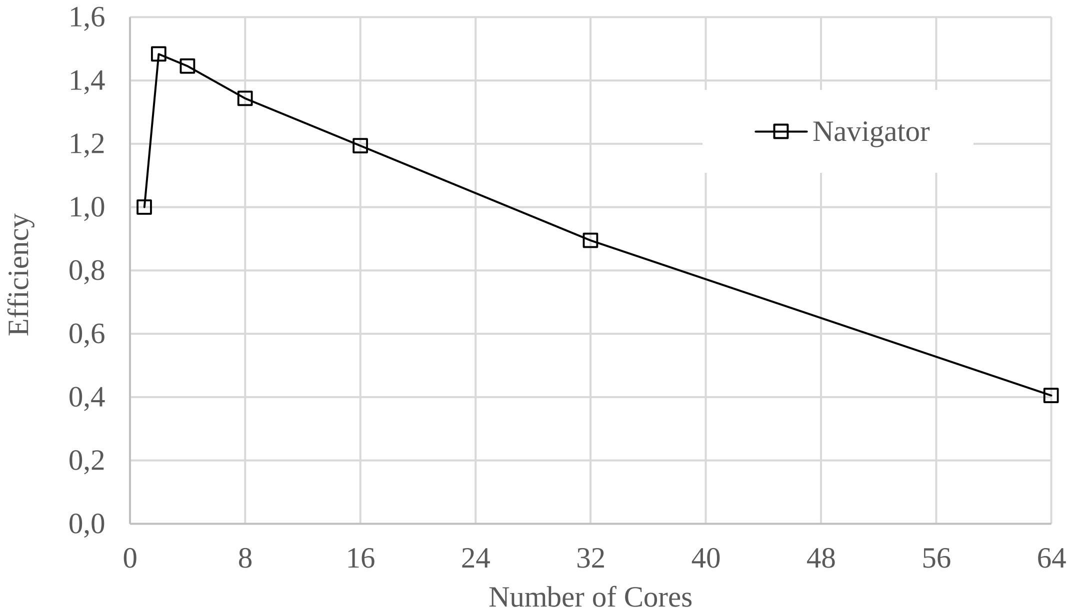

Figure 6 - Efficiency for a 1.7 million cell mesh, considering a single node.

Scalability and Parameter Tunning

To assess the influence of the model dimension, a new mesh is generated for the previously described problem. The new mesh is comprised of approximately 12.7 million cells. This mesh density corresponds to the one which is required to perform the final aero-acoustic numerical simulation.

A speedup and efficiency analyses were performed for the new mesh density and, in Figure 7 and Figure 8, the results obtained are compared with the ones from the 1.7 million cells mesh. In general, the shape of the speedup curve is similar for both cases, but the curve corresponding to the denser mesh presents a downwards offset. In addition, the superlinear speedup when moving from one to two cores is much less pronounced.

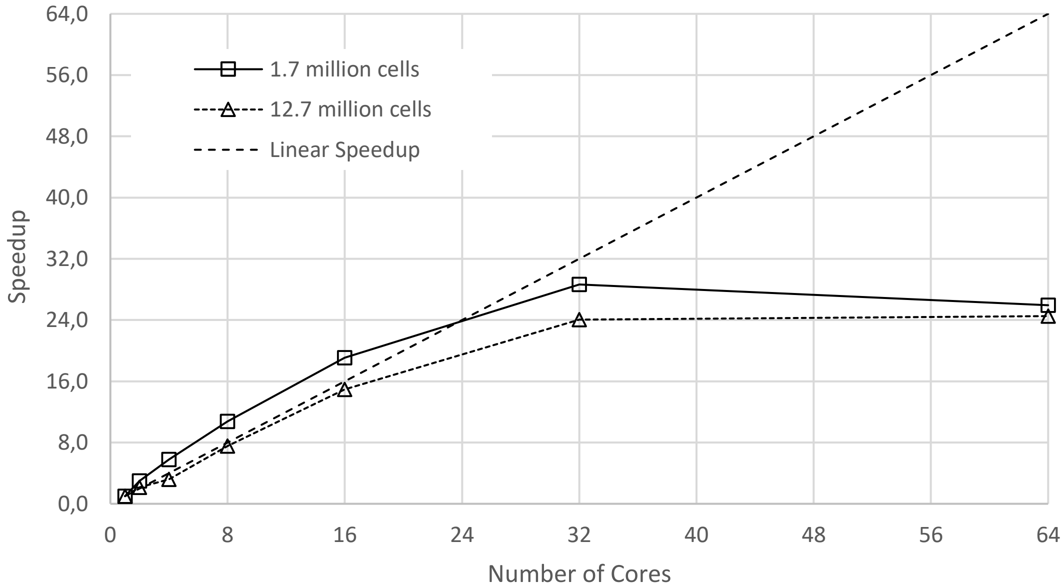

Figure 7 - Speedup for two mesh densities, considering a single node.

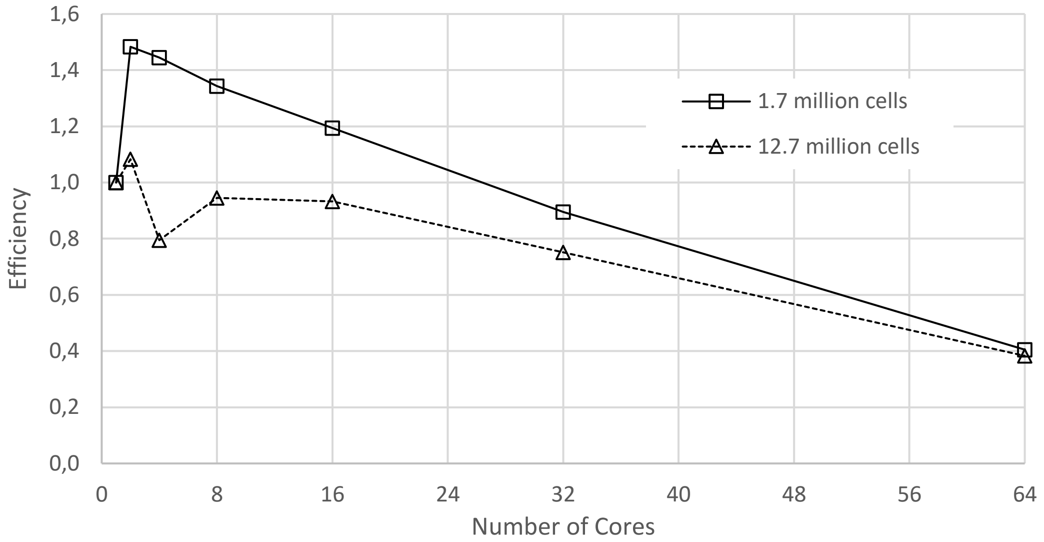

Figure 8 - Efficiency for two mesh densities, considering a single node.

Finally, to determine if it is possible to increase the performance of the Navigator cluster, the simulation with the finer mesh is executed considering the number of cores equally distributed along 2 nodes, instead of using one. On the Navigator cluster, a single node on the cpu2 partition has 40 cores available, so when requesting for more than 40 cores with no hyperthreading, two nodes are automatically selected. However, for this analysis, the number of cores in each node was forced by using the SLURM command '--ntasks-per-node' to have the same cpu load in both nodes. The results are given in Figure 9 and Figure 10. As can be seen, the results are quite similar up to 32 cores. However, when considering 64 cores there was an increase of about 13.4 and 0.21 in the speedup and efficiency, respectively. This demonstrates that, when considering a high number of cores, it is advantageous to divide them equally among each node. Comparing with the research workstation, the use of HPC allows to reduce the execution time by a factor of approximately 12.3.

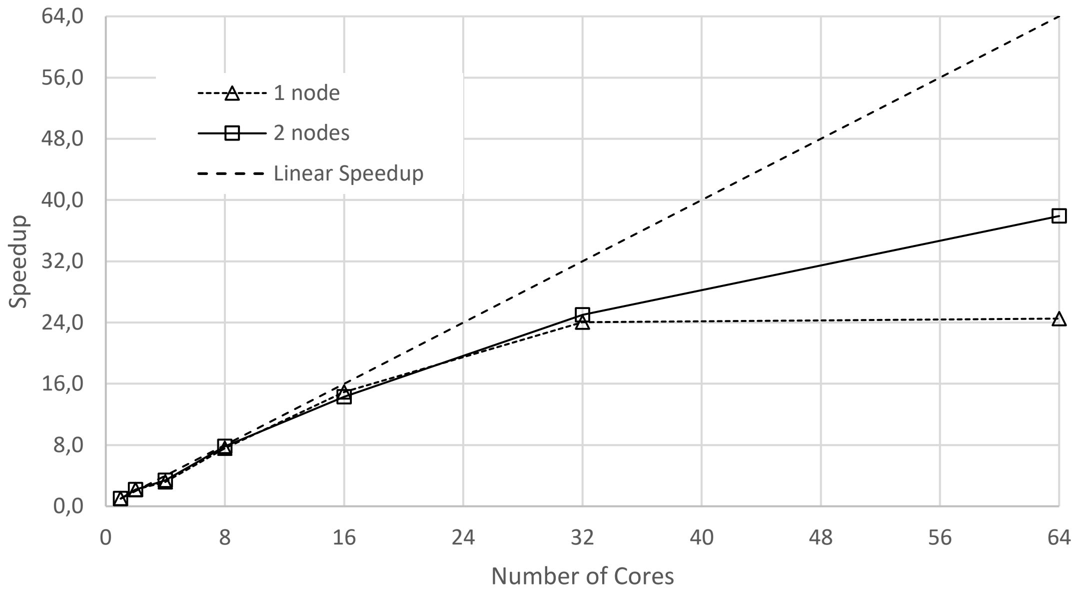

Figure 9 - Speedup using one and two nodes.

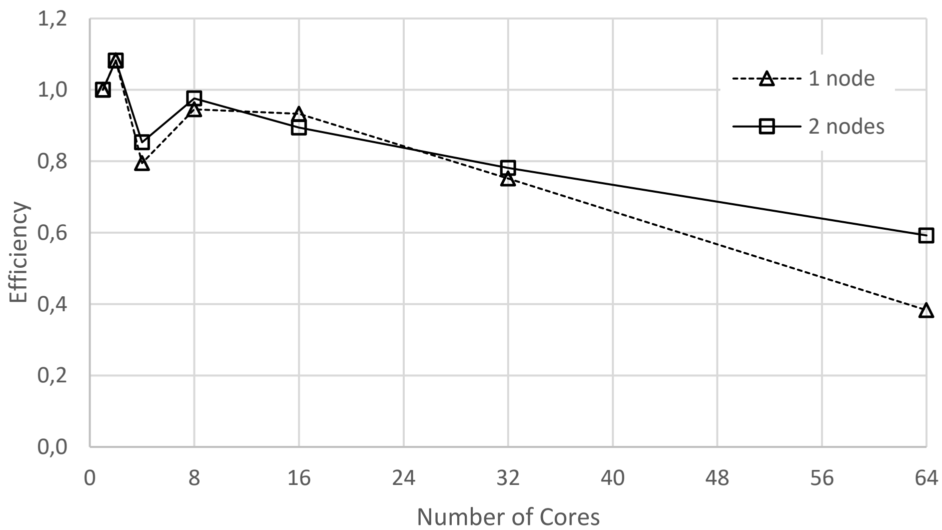

Figure 10 - Efficiency using one and two nodes.

Aero-Acoustic Simulation

In order to perform an aero-acoustic analysis, the simulation must be able to capture the flow unsteadiness arising from turbulence. Reynold-Average Navier-Stokes (RANS) models are, in general, not able to capture the required unsteadiness. Therefore, it is necessary to use Scale Resolving Simulations (SRS) that can capture the unsteady motion of a range of different turbulent scales. These methods are, however, very time-consuming since they require very fine meshes to be able to capture the turbulence effects of the flow. In this work, the Detached Eddy Simulation (DES) is considered.

The simulation process of the ventilation model can be divided into three main stages, which are described in the following:

- In the first stage, a steady-state RANS model is used in order to obtain a converged solution that acts as the starting point for the following step of the simulation.
- In the second stage, a transient simulation using the DES model is performed. The simulation must run for a sufficient number of increments in order to guarantee that the flow is fully developed. In addition, to capture the motions and to be able to perform the acoustic analysis in the following stage, the time step must be small. For the current model, a time step of 5x10-5 seconds (flow time) is considered. To obtain a fully developed flow, a flow time of 0.25 s must be achieved. This results in 5000 time steps and, at each time step 5-20 iterations are performed in order to obtain convergence at each time step. As can be seen, this results in a very high number of iterations.
- In the final stage, and since the flow is now fully developed, the acoustic model is turned on and the process is continued for an additional 0.1 to 0.15 flow time seconds. The acoustic data is written for each time step and can be post-processed after completion.

The first step of the analysis was already performed in the previous section using the 12.7 million element mesh. For the remaining calculations, a Navigator configuration using 64 cores divided equally into two nodes is considered, since it performed better in the previous analysis. The calculation of the second stage required approximately 96 hours (4 days) of execution time. The third stage required an additional 48h, summing up to a total of 144h.

If the current simulation was to be performed in the research workstation limited to 4 cores, it would require around 1771h (about 74 days) of execution time. This result clearly demonstrates the main advantage of using HPC resources for the numerical simulation of complex problems.

Figure 11 presents the comparison between the numerical and experimental data. The noise is measured at a distance of 500 mm from the front of the airvent. As can be seen, there is a good agreement with the experimental data, specially for the higher frequencies. There are some variations for the middle-range frequencies, so additional simulations may be required to fully calibrate the model. For instance, using different SRS such as Scale-Adaptive Simulation (SAS) or Large Eddy Simulation (LES) or extending the volume corresponding to the vehicle interior. However, since the main objective of the present project is to evaluate the gain of using HPC for an industrial complex problem, the calibration of the numerical model is outside the scope of the work. This is, nevertheless, the first step to obtain the calibrated model.

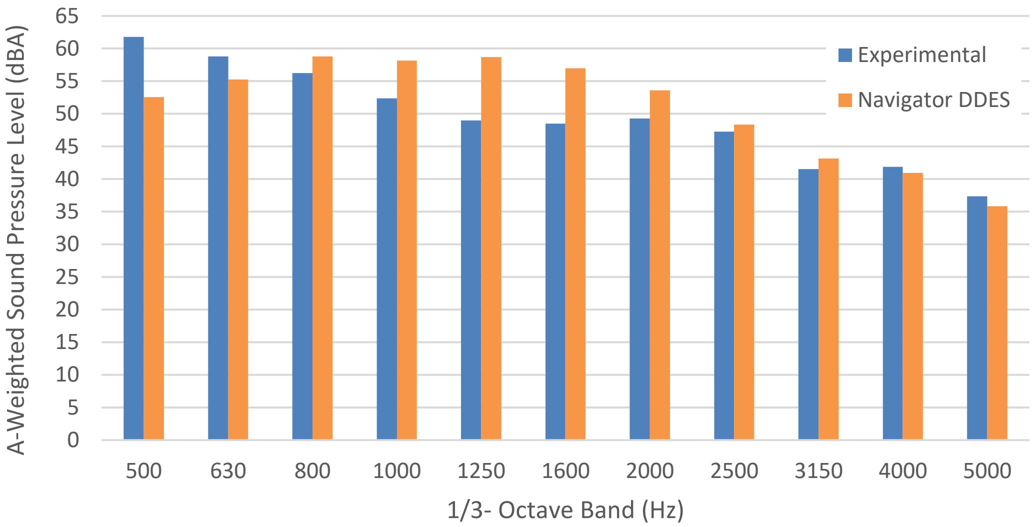

Figure 11 - Comparison between numerical and experimental data.

The velocity profile for the steady-state solution (flow time of 0s) and for a flow time of 0.25s are given in Figure 12. As can be seen, the steady-state case, which employs a Reynolds-Average Navier-Stokes approach, is only able to represent the behaviour of the mean flow. On the other hand, the model using the DES approach is able to account for the effects of turbulence in more detail and this is a crucial factor to be able to perform an accurate aero-acoustic simulation. However, this approach requires significantly higher computational resources.

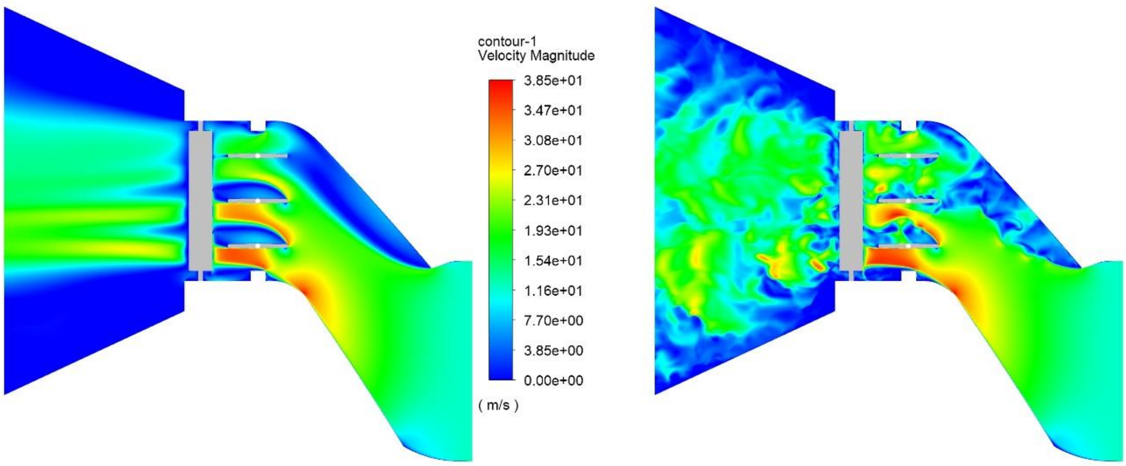

Figure 12 - Flow velocity for (left) steady-state RANS and (right) flow time 0.25s with DES model.

Conclusions

In this work, the performance of a HPC cluster was assessed when solving a complex Computational Fluid Dynamics problem. The model corresponds to a simplified automotive airvent. The goal was to perform an aero-acoustic analysis to determine the noise generated due to airflow and turbulence effects and compare the results with experimental data available. This kind of problems require significant computational effort since fine meshes and small time steps are necessary.

In a first stage, the performance of the Navigator cluster was assessed, followed by some parameter tunning to determine a more efficient distribution of the computational resources. Is was shown that there is speedup of about 28 when considering 32 cores on a single node for a mesh of about 1.7 million cells. This also represents a decrease of the execution time by a factor of 7when compared to a research workstation limited (by the simulation software) to 4 cores. A superlinear speedup was also visible when moving from one to two computation cores.

Increasing the problem size to 12.7 million cells, the superlinear speedup was less pronounced, but the overall speedup and efficiency rates were similar. Since each node in the cluster contains 40 cores, it was seen that, for 64 cores, when the computational load was distributed equally between the two nodes (i.e. 32 cores per node) the speedup and efficiency increased by 13.4 and 0.21, respectively. When compared to the research workstation running on 4 cores, the use of HPC allowed to reduce the execution time by a factor of approximately 12.3.

Finally, the aero-acoustic simulation of the airvent was performed and the obtained results are in good agreement with the available experimental data for higher frequencies. The complete numerical simulation required a total of about 144h (6 days) using 64 cores on two nodes. The same problem would take approximately 1771h (74 days) if solved in the research workstation using 4 cores. Therefore, it can be seen that there are significant benefits when considering High-Performance Computing. This result is especially relevant in the context of Small or Medium Enterprises (SMEs) that do not have available in-house the processing hardware required to solve complex and time consuming computational models.

The numerical simulation of the airflow in an automotive airvent has significant interest for the mouldmaking and plastic industries since it allows to develop new concepts and designs at an early project stage, allowing to reduce or avoid modifications of the product at later stages that would imply greater costs. In the case of Portugal, the mouldmaking and plastic industries are composed of a high number of SMEs that are not able to have the necessary hardware to solve such problems in a reasonable amount of time. These companies are the ones that should be able have higher benefits when using HPC, by optimizing their processes, increasing productivity and reducing costs related to new product development. The case study presented in this work is only one of the possible applications of HPC in the industry. With concepts such as Big Data, Internet of Things (IoT) and Digital Twin, the use of HPC can give the necessary competitive edge to SMEs.

As future works, it is intended to perform a deeper calibration of the numerical model, followed by a numerical and experimental testing of a real case airvent. The final goal is to develop a numerical tool that can be used to assess the acoustic performance of different designs in the early project stages of the development of a new product.

References

[1] P. A. Jennings, G. Dunne, R. Williams e S. Giudice, Tools and techniques for understanding the fundamentals of automotive sound quality, Proceedings of the Institution of Mechanical Engineers, Part D: Journal of Automobile Engineering, 224: 1263–1278, 2010.

[2] Ansys® Fluent, Release 2020 R2, ANSYS, Inc.

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 823767. The authors would also like to acknowledge Prof. A. Gameiro (University of Coimbra) and Dr. M. Fontul (Iber-Oleff) for sharing their expertise and for the scientific discussions related to the current project. The support of UC-LCA administration team, in particular Bruno Andrade, for setting up and testing ANSYS Fluent for the Navigator cluster environment is also acknowledged.