Available online at www.prace-ri.eu

Partnership for Advanced Computing in Europe

Enhanced Airflow Simulations around Filling Machines in Clean Rooms

R. Eisenschmid[a]*, B. Große-Wöhrmann[b]

[a]*OPTIMA pharma GmbH, Otto-Hahn-Strasse 1, 74523 Schwäbisch Hall, Germany*
[b]*High Performance Computing Center (HLRS), Nobelstr. 19, 70569 Stuttgart, Germany*

Abstract

In this white paper, the main results of the SHAPE pilot project "Enhanced airflow simulations around filling machines in clean rooms" are presented. During this project on the Tier-0 system Cray XE6 "Hermit" at HLRS, Germany, the open source CFD software package OpenFOAM® v2.2.2 was utilized to run simulations meeting the requirements of industrial production. Besides testing different turbulence models, emphasis was placed on parallel mesh generation with *snappyHexMesh* and the decomposition and reconstruction of large meshes with more than 10 million cells.

Furthermore, the cooperation between PRACE and the SME OPTIMA pharma GmbH is described, and a report on the benefits for the SME, the lessons learned and the future activities is added.

1. Introduction

OPTIMA pharma GmbH [1], located at Schwäbisch Hall, Germany, produces and develops filling and packaging machines for pharmaceutical products - sterile and non-sterile liquids - and pharmaceutical freeze-drying systems as well as isolator (clean room) and containment technology. The company is operating worldwide with 600 employees; it constitutes the pharma division of OPTIMA packaging group GmbH with 1800 employees worldwide.

Sterile filling lines are enclosed in clean rooms, and a detailed and reliable knowledge of the airflow inside the clean rooms would support the Computer Aided Engineering (CAE) process and enhance the design of the filling machines according to the customers' requirements. For example, the airflow around the small bottles, called vails, and cartridges in the filling line should prevent the remaining dust and impurities from reaching the sterile materials. In addition, as the pharmaceutical products may be toxic, the clean rooms have to be vented and purged, and the contaminated, outgoing air should leave the clean rooms through the ventilation slots without endangering the staff. Furthermore, turbulences and flow detachments, especially in corners and on filling devices, should be avoided.

OPTIMA pharma GmbH was chosen as one of ten projects in the PRACE SHAPE (SME HPC Adoption Programme in Europe) pilot [2]. The SHAPE programme is designed to equip European Small and Medium-sized Enterprises (SMEs) with the awareness and expertise necessary to take advantage of the innovation possibilities opened up by HPC, thus increasing their competitiveness. With the support of HPC experts, the SMEs develop HPC solutions relevant for their investigation and production processes.

---

* Corresponding author. *E-mail address*: Ralph.Eisenschmid@optima-pharma.com

Figure 1: Isolation technology produced by OPTIMA pharma.

The goal of this SHAPE pilot project was to simulate the airflow in a clean room with OpenFOAM® on a Tier-0 system meeting the requirements of industrial production. We were granted a PRACE preparatory access type C project on the Cray XE6 "Hermit" at HLRS, Germany, and utilized the version OpenFOAM® v2.2.2 [3] already installed on Hermit to run the tests and simulations. OpenFOAM® is an open source CFD software package offering a variety of turbulence models as well as serial and parallel tools for mesh generation and mesh decomposition. Information about wall-clock time and memory used was collected by the Cray Profiler [4].

The project work included the development and testing of reliable Linux scripts for efficient long-term usage: One objective was the efficient and error-free adjustment of the Cray batch scripts and OpenFOAM®'s dictionaries to, for example, varying core numbers and different use cases; another challenge was how to extract the relevant pieces of information from the Cray Profiler output files in order to easily import them into Excel.

The main topics were parallel mesh generation (see section 2), domain decomposition and mesh reconstruction (section 3) and turbulence modelling (section 4). Furthermore, we report on our experiences and lessons learned during this SHAPE pilot project in section 5.

2. Parallel mesh generation with *snappyHexMesh*

The parallel tool *snappyHexMesh* [5] generates meshes snapped to surfaces defined, for example, by CAD data in STL format. It offers the possibility to vary the number of refinement levels (r) of the initial block mesh as well as the number of cells between refinement levels (c) and the number of added surface layers (s). We were interested in how these three parameters influence the final number of mesh cells, the wall time and the memory used by *snappyHexMesh*. We studied a very simple geometry, a cube (see Figure 2 and subsection 2.1), and a complex geometry: the mock-up of a RABS (Restricted Area Barrier System, which is a clean room), see Figure 4 and subsection 2.2.

2.1. The cube

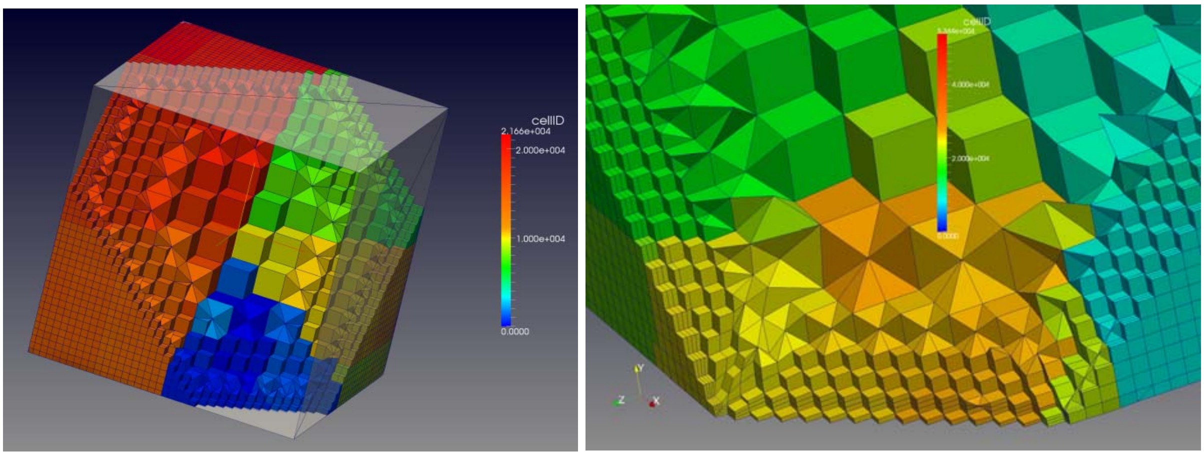

Figure 2: Left: Cube with r=2 refinement levels and c=3 cells between levels. Right: Cube with r=2 and s=4 surface layers.

Table 1: Measured wall-clock time, memory used and final number of mesh cells for c=3 cells between refinement levels, s=0 surface layers and the number of refinement levels r varying from 2 to 7.

refinement level *r*	**wall-clock time in sec**	**memory used in MB * processor**	**number of mesh cells**
2	400	37.8 * 64	21664
3	650	60.7 * 64	117776
4	1230	87.7 * 64	662915
5	2520	114.7 * 64	2261944
6	9830	327.3 * 64	10413080
7	41460	1097.9 * 64	34898928

Table 1 shows the measured wall-clock time, memory used and final number of mesh cells for the cube, depending only on the parameter r. Our goal was to find an approximate formula fitting the number of mesh cells $n(r)$. On each refinement level, the cells near surfaces are split in $2^3 = 8$ sub cells. Therefore, we make the following approach for the total number n of mesh cells depending on the number of refinement levels r:

$$n(r) = \left(z^0 + z^1 + z^2 + \ldots + z^r\right)\frac{A}{d^2} + \frac{V}{d^3} - \frac{A}{d^2}.$$

Here, z denotes a real number between 4 and 8, V is the volume, A is the size of the surface, and d is the original cell size. Therefore, A/d^2 is approximately the number of mesh cells at the surface, and $V/d^3 - A/d^2$ is the number of mesh cells in the interior. Summing the terms up, we get the simple approximation

$$n(r) = \frac{z^{r+1} - 1}{z - 1}\frac{A}{d^2} + \frac{V}{d^3} - \frac{A}{d^2}.$$

Figure 3 shows that z=4.6 and z=4.8 are well approximating $n(r)$ for large values of r, whereas for small values of r, z=5.0 is the better guess.

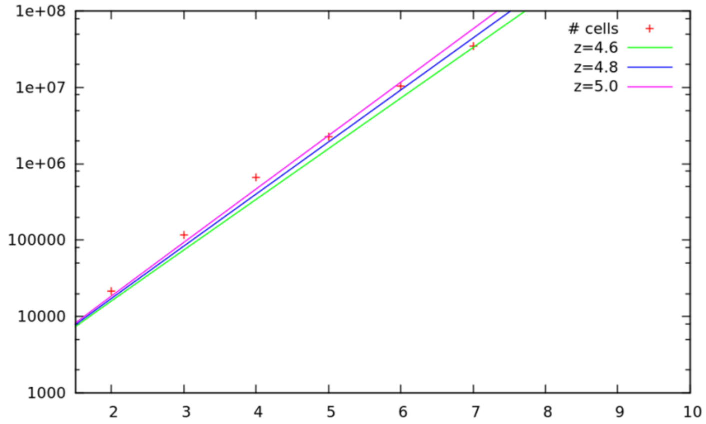

Figure 3: Fitting $n(r)$. The horizontal axis is the refinement level r, and the vertical axis is the cell number $n(r)$.

2.2. The mock-up

Mock-up models are typically physical 1:1 constructions made of wallboard with reduced details. Main purpose of these models is, to encourage people to test handling with the (wooden) devices of the later clean room in an early stage of design and development. The mock-up CAD models can be used also preferred to set up a CFD simulation, because the mock-up CAD data are much smaller than the full CAD design.

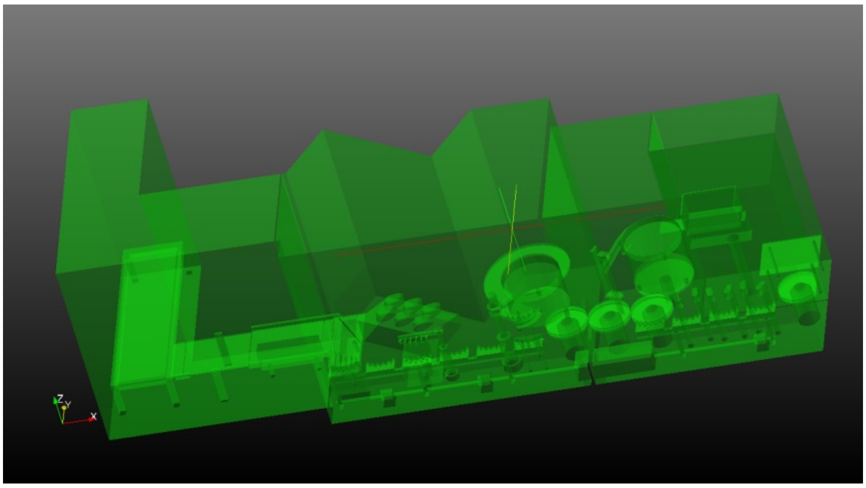

Figure 4: Clean Room: Mock-Up model of a RABS with aseptic liquid filling machines inside (CAD data).

Figure 4 presents the complex geometry of a filling line in a clean room. Filling devices, hoppers, sorter devices and transport chain are displayed.

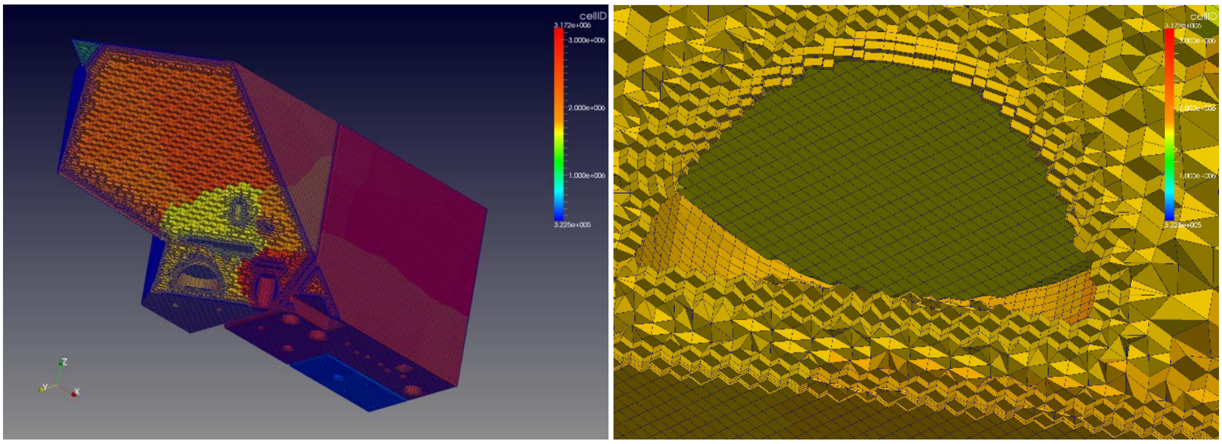

Figure 5: Clean room mesh example with r=2 and s=4. Right: detail.

Figure 5 shows an automatically generated mesh by the OpenFOAM tool *snappyHexMesh* with refinement level 2 and with 4 surface layers (see detail image) on all boundaries. Mesh quality and surface layer coverage is quite satisfying, just some small details have some missing surface layers.

3. Domain decomposition and mesh reconstruction

The serial tools *decomposePar*, *reconstructParMesh* and *reconstructPar* are commonly used for mesh decomposition after running *snappyHexMesh* and for reconstruction the mesh and the fields after the solver run. They are major bottlenecks: We found insufficient performance of these serial tools on large meshes with more than 40 million (40 M) cells. Furthermore, they require more memory than usually available: *decomposePar* runs out of memory on an ordinary 64 GB node when decomposing snapped meshes with more than 40 M cells. Similarly, *reconstructParMesh* can handle meshes only up to 30 to 40 M cells on a 64 GB node. In addition, *reconstructPar* runs out of memory for meshes with more than 30 M cells, and more seriously, it runs out of the maximal wall-clock time (24 h) on Hermit on even smaller meshes if there are many time steps (write-out times).

A possible workaround of the memory problem would have been the utilization of special pre- and post-processing nodes with 1 TB memory available on Hermit, but we found another solution and successfully applied it: The usage of the parallel OpenFOAM® tools *renumberMesh* and *potentialFoam.* This means that the serial processes can partly be substituted by parallel processes. In this way, meshes with up to 512 M cells were performed and the reduction of wall-clock time was respectable: The parallel tools were in the region of 80%-90% faster.

4. Turbulence modelling and scaling tests

We took a nearly direct numerical simulation (DNS) of the Karman vortex street as benchmark for the turbulence model (see Figure 6). The Karman vortex street seems to be the most suitable case for transient flow simulation benchmarks because it is well investigated, and a simple formula for the vortex frequency is also well-known (see below). DNS is the nearest approximation to reality.

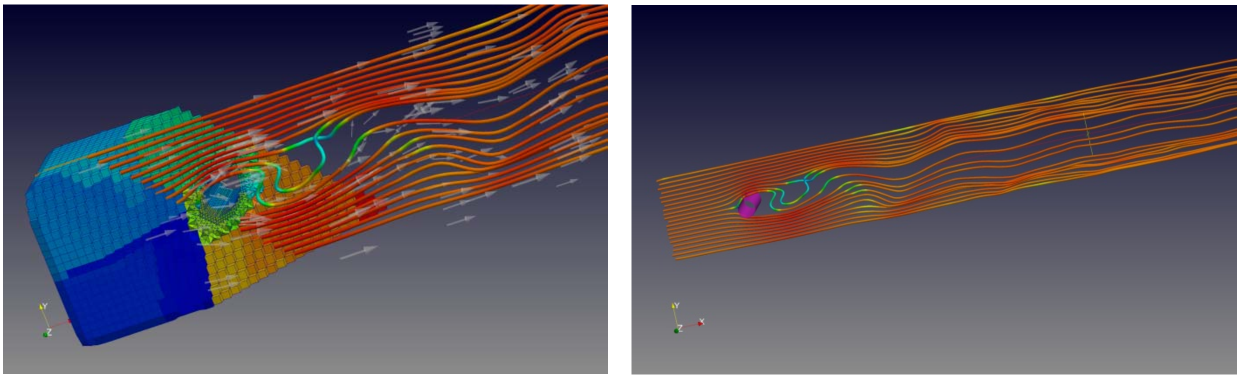

Figure 6: Karman vortex street.

The Karman simulation shows good compliance with the theory of the Karman vortex street. With frequency f, Strouhal Number Sr, velocity U=0.45 m/s and diameter d=0.2 m of the cylinder, one gets

$$f = \frac{Sr * U}{d} = \frac{0.2 * 0.45\,\mathrm{m/s}}{0.2\,\mathrm{m}} = 0.45\,\mathrm{Hz}$$

Post-processing shows a wavelength of about 1.0 m, as expected. Therefore, this Karman-DNS case can definitely be used as benchmark for transient turbulence modellings like Delayed Detached Eddy Simulation (DDES) and Improved Delayed Detached Eddy Simulation (IDDES). Figure 7 shows first results of the Karman case simulated with DDES.

4.1. Spalart-Allmaras (I)DDES

We set up the Spalart-Allmaras DDES and IDDES cases with OpenFOAM's solvers *pisoFoam* and *pimpleFoam* following [6] and [7]. The boundary conditions for Spalart-Allmaras (I)DDES were chosen according to the approximations given in section 6.5 of [6]:

$$k_{\mathrm{in}} \approx 1.5\left(u_{\mathrm{in}} t_{\mathrm{in}}\right)^2$$

$$\nu_{T,\mathrm{in}} \approx \sqrt{1.5} u_{\mathrm{in}} t_{\mathrm{in}} L_{T,\mathrm{in}}$$

where u_{in} is the known mean velocity at the inlet boundary, t_{in} is the turbulence intensity at the inlet, and the length $L_{T,\mathrm{in}}$ of the largest turbulent elements can be approximated by the hydraulic diameter $D_{H,\mathrm{in}}$ of the inlet aperture: $L_{T,\mathrm{in}} \approx 0.1\ D_{H,\mathrm{in}}$. With u_{in}=0.45 m/s, t_{in}=0.01and $D_{H,\mathrm{in}}$=1 m we got $k_{\mathrm{in}}=3*10^{-5}$ $\mathrm{m}^2/\mathrm{s}^2$ as inlet boundary condition for the turbulent kinetic energy and $\nu_{T,\mathrm{in}}=5.5*10^{-4}$ m^2/s for the turbulent kinematic viscosity.

As we did not utilize wall functions to model the flow near walls, we needed a finer mesh resolution with surface layers near the walls. Needed is a dimensionless wall distance $y^+ \leq 2$. OpenFOAM's tool *yPlusLES* gives values between 0.7 and 4.3 which is the right order of magnitude but indicates that the mesh near surfaces and corners has to be refined and enhanced.

Calculation of the wall distance y for the Karman cylinder (see [8]):

U=0.45 m/s, L=0.3 m (half contour) and viscosity $\nu=1.5*10^{-5}\mathrm{m}^2$/s give Reynolds number Re=8900. y^+=30 fits y=1.5 cm (cell thickness normal to wall: 3 cm), and y^+=2 fits y=1mm (cell thickness normal to wall: 2 mm).

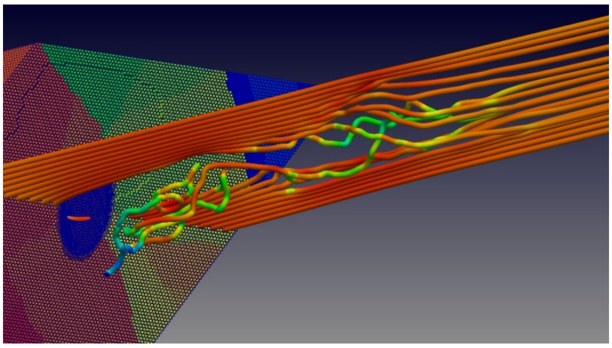

Figure 7: The Karman case simulated with DDES. Left: flow tubes. Right: pressure iso-surfaces.

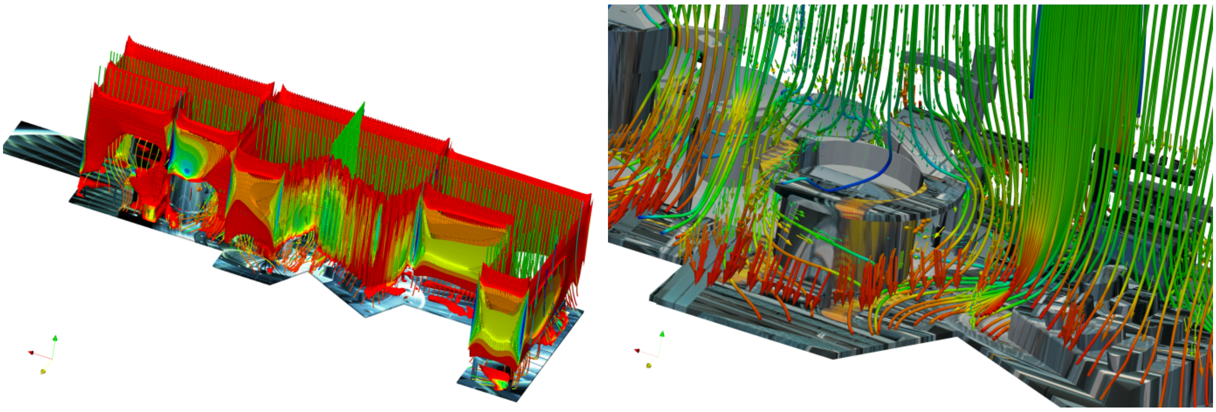

Figure 8: Cleanroom Mock-Up case simulated with stationary RANS modelled with the *simpleFoam* solver. Left: streamlines and k- iso-surfaces. Right: detail with streamlines

A stationary adaption of air flow in clean rooms is shown in Figure 8. The simulation was run with the *simpleFoam* k-omega-SST solver. Detached eddies and other typical transient effects can't be resolved by this stationary solutions. Another issue are the cellsizes of the nearest cells to the surface (surface contact). As shown in RANS theory, center of minimum cellsize should be $y^+ > 30$. For this CFD case, wall-function surface cells should have 30mm thickness! So it is impossible to resolve CAD structures below 60mm. The Reynolds Number Re is 9400, after all. This seems to be too small for RANS models. Maybe minimum Re >1e5 would help to get a case, which runs better for RANS models. A benefit of stationary solvers like *simpleFoam* is, that they run very fast compared to transient solvers like DDES.

4.2. Scaling Tests

OpenFOAM toolboxes show a very bad scalability. Good strong scaling only could be performed up to 128 cores (e.g. *snappyHexMesh*) or 256 cores (e.g. *simpleFoam*). The only reasonable use of (weak) scaling is to utilize memory between 50% and 80% load, then OpenFOAM runs quite effectively (see Table 2 and Table 3). The performance of OpenFOAM for these models was disappointing. The reasons for the poor performance were extensive I/O operations, where large numbers of files are written on file systems, and also processor/processor communication and non-optimized OpenFOAM codes were bottlenecks.

In particular, OpenFOAM's extensive I/O operations were most problematic. Initially there was an issue where the computing time ran out and processes ended without any error messages of the profiler or the queue or Tier-0 system executables. No results were written on the file system, but wall-clock time ran out. This was a source of some frustration to the team, however after some investigation it was discovered that the Tier-0 system's disk quota was being exceeded due to the number of files being generated by OpenFOAM. Every process generates about 50 files (dependent on the number of written time steps) so that massively parallel jobs with 4096 cores generate about 200000 files. This number is 20% of the disk quota of 1 million files! This is the reason why only a small number of scaling tests could be performed.

Only some *snappyHexMesh* scaling jobs ran without quota problems. The results can be seen in following tables.

Table 2: Strong Scaling Tests

Number of cores	Number of cells	Wall-clock time	Speed-up vs the first one	Number of nodes	Number of processes	Speed-up ideal	Speed-up efficiency
64	19937000	180	1	2	64	1	1
256	19937000	112	1.61	8	256	4	0.40
512	19937000	188	0.96	16	512	8	0.12
1024	19937000	802	0.22	32	1024	16	0.01
4096	19937000	fault	fault	128	4096	64	fault

Table 3: Scaling Tests with constant core number for estimation of the maximum possible cell number

Number of cores	Number of cells	Wall clock time [s]	Speed-up vs the first one	Number of nodes	Number of process	Speed-up efficiency	Approx. number of cells/core	efficiency cells per core and [s]
1024	19937000	802	1	32	1024	1	19470	24.28
1024	25672548	776	1.03	32	1024	1.03	25071	32.31
1024	48424264	852	0.94	32	1024	0.94	47290	55.50
1024	90582324	1124	0.71	32	1024	0.71	88460	78.70
1024	143314458	4001	0.20	32	1024	0.20	139956	34.98
1024	196049726	4074	0.20	32	1024	0.20	191455	46.99
1024	512348684	4275	0.19	32	1024	0.19	500341	117.04

5. More about the SHAPE pilot activity

In this chapter we provide further information about the SHAPE pilot, the collaboration between PRACE and OPTIMA, our experiences and future activities. The SHAPE team members were Ralph Eisenschmid, Bärbel Große-Wöhrmann (SHAPE coach) and Martin Winter (HLRS, CFD expert, internal consultant).

5.1. PRACE cooperation

In close collaboration with the other two SHAPE team members, the industrial partner R. Eisenschmid ran the OpenFOAM® test cases on the machine Hermit within the granted PRACE preparatory access type C project 2010PA2080. There were frequent telephone calls, often daily, and weekly or bi-weekly meetings at the HLRS.

In addition, PRACE facilitated R. Eisenschmid to participate in the PRACEDays14 in Barcelona and to present the project.

5.2. Benefits for the SME

OPTIMA had gained some experience with Tier-1 HPC systems at HLRS before the SHAPE pilot project. Nevertheless, without the help and continuous support of the SHAPE coach (B. Große-Wöhrmann), OPTIMA would not have been able to set up and run simulations on the Hermit Tier-0 system. PRACE SHAPE support began with writing the PRACE preparatory access proposal and continued with instrumenting the Cray Profiler, editing the scripts and setting the turbulence models up. Additionally, the SHAPE coach fixed problems that occurred while porting scripts from OpenFOAM version 2.1.1 used on the Tier-1 system to version 2.2.2 used on Hermit.

This SHAPE project enabled OPTIMA to utilize OpenFOAM on the Tier-0 system Hermit for their industrial development and design processes. With the findings summarized in this paper, they are now able to set up a CFD case in considerably reduced time (about 80% savings). Prediction of mesh sizes, processor resources and wall-clock time of all OpenFOAM processes helps to optimize the HPC case and to save much money and time. Therefore, OPTIMA can use full HPC capacity with a minimal waste of resources and very reduced queuing times (jobs with runtimes predicted as less than 4 hours). The results will help to select the most appropriate solver for handling air flows in clean rooms at a maximum of accuracy and a minimum of resources. See also section 5.4 about future activities.

Finally, SHAPE provides good opportunities to contact other OpenFOAM users at events like PRACEDays14 and to present poster images on the PRACE booth at the ISC14.

5.3. Lessons learned

Generally speaking, the concept of the SHAPE pilot worked very well: supporting the industrial partner in writing the PRACE preparatory access application and in becoming acquainted with the possibilities provided by a Tier-0 system was sensible and successful. The industrial partner was quite surprised that the work on a Tier-0 system appeared not as predictable as expected. The availability of HPC systems is not about 99% like usually required in pharma systems. Furthermore, the queuing time strongly depends on the workload of other users. Another time consuming issue is the occasional overloading of the front nodes, if users run commands that demand more resources than intended. Therefore, more single processor shared front nodes appear desirable.

Another lesson learned is that the work took more time than foreseen. The industrial partner was so intent upon preparing the OpenFOAM cases for efficient, long-term usage that the time was quite short for this SHAPE project: there were only three months between the start of the preparatory access at the end of January and the writing of the final documents beginning in May. Therefore, future SHAPE projects should be better coordinated with the allocation periods of the PRACE preparatory access projects. Furthermore, it cannot be assumed that the partners from industry have no other duties besides their SHAPE projects. Altogether, the work has taken much more time than two personal months up to now, and the preparatory access project is still running till the beginning of August. So PRACE members and also industrial partners like Optima have to learn these lessons. To summarise: industrial partners expectations need to be managed, for example if something goes wrong with the HPC system, or the PRACE support takes longer as expected, or timelines are exceeded; similarly, PRACE members have to understand that typically industrial partners cannot spend 100% of their working time on PRACE projects.

5.4. Future activities

During this SHAPE pilot we focused our attention on the question whether better chosen turbulence models and finer meshes will produce qualitatively better results revealing unknown details and providing deeper insight into the flow. We adequately integrated the new results and methods in our design processes. With a better knowledge of the airflow in our clean rooms, we are now able to enhance the design of our filling machines according to our customers' requirements. This optimized set of parameter files and OpenFOAM dictionaries will be used as template for many other clean room geometries and airflow cases. Solver independent findings of this project, like optimized meshing and decomposition, can be implemented to solve new problems and tasks like multiphase studies with OpenFOAM's *interFoam* solver, for example. In the future, airflow simulations will help to replace most of the expensive smoke studies in production machines at the customer's site. Simulations will also attend the design process in an early state and support the whole CAE process. Most of the negative developments can be avoided at an early stage, and trial and error cycles and expensive redesigns and reconstructions can be minimized. OPTIMA will continue the successful collaboration with HLRS as an industrial customer. Furthermore, OpenFOAM is widely used in both academia and industry, and the results of this project may help other OpenFOAM users to set up their cases and to run their applications.

6. Conclusions

It was a very interesting and challenging project. We were glad to get the opportunity to participate in the SHAPE pilot.

References

[1] www.optima-pharma.com
[2] www.prace-ri.eu/shape
[3] www.openfoam.org
[4] wickie.hlrs.de/platforms/index.php/CRAY_XE6_Cray_Tools
[5] www.openfoam.org/docs/user/mesh.php
[6] Rüdiger Schwarze, *CFD-Modellierung, Grundlagen und Anwendungen bei Strömungsprozessen,* Springer-Verlag 2013.
[7] Martin Gramlich, *Numerical Investigations of the Unsteady Flow in the Stuttgart Swirl Generator with OpenFOAM,* Master's thesis in Fluid Dynamics, Chalmers University of Technology, Gothenburg, Sweden, 2012.
[8] www.cfd-online.com/Wiki/Main_Page

Acknowledgements

This work was financially supported by the PRACE project funded in part by the EUs 7th Framework Programme (FP7/2007-2013) under grant agreement no. RI-312763.

We acknowledge PRACE for awarding us access to resource Hermit based in Germany at HLRS. The support of Bärbel Große-Wöhrmann from HLRS, Germany, to the technical work is gratefully acknowledged.