OpenFOAM performance and scalability on various HPC systems

L. Koloszár, R. Giammanco, Ph. Planquart

von Karman Institute, Ch. de Waterloo 72. B-1640, Rhode-St-Genese, Belgium, +3223599763, koloszar@vki.ac.be

INTRODUCTION

The OpenFOAM simulation environment allows parallel computation in the field of fluid dynamics without additional license cost due to its open source nature. As such, efficient parallelization is of high interest since the number of processor can be used to perform a given numerical simulation is restricted only by the available infrastructure.

The current study examines different workstations and HPC systems to determine the effect of the underlying hardware using OpenFOAM 4.0. The focus is given not just the overall performance of the different infrastructures, but to their scaling.

Moreover, we were examining three different complexity of numerical setup: a single phase turbulent problem of natural convection flow (LES of 48M cells), a multiphase case with VOF method (laminar of 38M cells) and a multiphase problem including conjugate heat transfer (URANS with 28M cells). We have compared the scaling of the three cases, investigating the consequence of the different modeling approaches used.

THE INVESTIGATED TEST CASES

We have chosen test cases that are representing prototype problems currently the most challenging for us at the von Karman Institute. They possess different complexity from a physical point of view, while always correspond to a large enough problem in terms of mesh size, such that intensive scaling study is possible to accomplish on them.

The first test case investigated was a single-phase natural convection boundary layer simulation. This unsteady problem was using Large-Eddy Simulation to account for the non-resolved turbulent structures. In order to account for heat transfer and density variation induced buoyancy, the *buoyantBussinesqPimpleFoam* solver was used. The mesh consisted of 48M cells, representing the largest case considered in this study.

The second problem aims to predict the liquid film dragged by a wire pulled out from a liquid metal bath. Since the velocity of the wire is still relatively slow, the flow is expected to stay laminar. Therefore, no modeling of turbulence is applied. To resolve the interface between the two phases the Volume of Fluid method is used, and the simulation was performed by the *interFoam* solver.

The third case is dealing with a multiphase system, as well, of a prototype nuclear reactor experimental facility. The primary coolant is liquid metal, while the containment is filled with Argon, as cover gas. Due to the high temperature gradients present in the system conjugate heat transfer in the solids are considered. The flow is turbulent in this case and the URANS approach is used to resolve it. The interface between the two phases again modeled with the Volume of Fluid Method. No native OpenFOAM solver could resolve all the physics involved, therefore, for this test case our own developed *myrrhaFoam* solver was used.

THE TESTED INFRASTRUCTURES

Workstations:	**Intel Genuine**	**AMD Ryzen**	**AMD Epyc**
Max. number of Threads	16	32	64
Location	VKI	VKI	VKI
Specs	2x E5-2630 CPUs X9DAi MB 16x8GBDDR3 1333MHz ECC RAM SMT ON	1x TR 1950X CPU X399 Taichi MB 8x16GB DDR4 2134 MHz NO ECC RAM SMT ON	2xEpic 7551 H11DSi-NT MB 16x32GB DDR4 2667 MHz ECC RAM SMT (ON?)

Clusters:	**CV**	**Thinking**	**Genius**	**Marconi A1**
Max. number of processors used	512	256	256	2304
Location	VKI	VSC	VSC	CINECA
Specs/Blade	4xOpteron 6376 H8DGU MB 32x 8GB DDR3 1600Mhz ECC RAM	10x Xeon E5- 2680v2 64GB 2800 MHz RAM	2x18 Intel Xeon 6140 192GB 2300 MHz RAM	2x18 Intel Xeon E5-2697 v4 128GB 2300 MHz RAM

EFFECT OF UNDERLYING HARDWARE

The second test case, simulation of wire withdrawal by Large-Eddy Simulation, was used to test the different hardware and infrastructures, in details. The simulation was initialized with the same initial condition for all the cases such that the same 40 time-steps were used to evaluate the performance.

Within these 40 time-steps the simulation data was exported 4 times in order to include the I/O time to the assessment.

The results obtained are summarized in Figure 1. All the simulation times were normalized by the time took to finish the simulation on the AMD Authentic cluster located at VKI. As it is expected, the absolute performance of the various architectures follows their release dates. Though, it is interesting to see that the Intel Skylake cluster located at VSC performs slightly worse than its predecessor, the Intel Broadwell. Though, the two architectures are not located at the same Institute, so the variation might be due to other factors than the performance of the processors, only. The new AMD Epyc processor has a very remarkable performance, though we had not the chance to test an Intel processor issued in the same year.

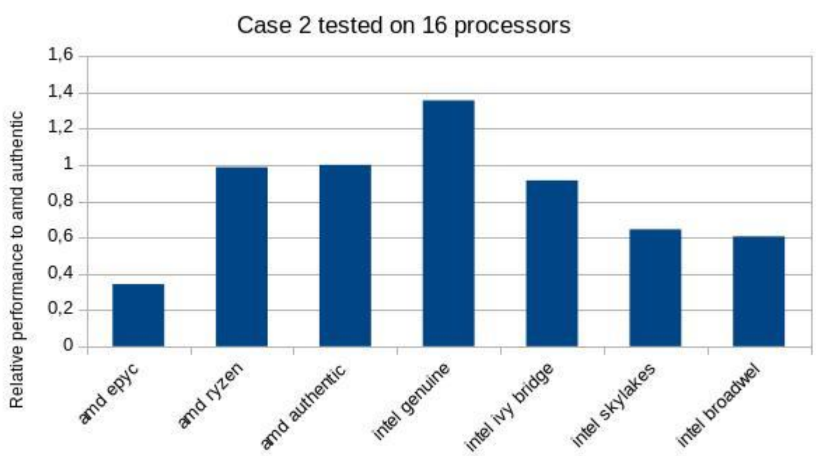

Figure 1: Performance of the different tested architectures: scaled by the performance of amd authentic.

A detailed scaling test will be shown in the presentation, but here in the abstract it is omitted for brevity.

EFFECT OF INVOLVED PHYSICAL MODELING

We have tested three different test cases to investigate how the scaling is changing if the underlying modeling is altered. We have moved from a single phase simulation to a multiphase with interface tracking. Finally, we introduced conjugate heat transfer to the multiphase system. The three cases were simulated with various numbers of processors on the CINECA cluster in the framework of the PRACE Preparatory Access project allowing to perform the scaling up to 2304 processors.

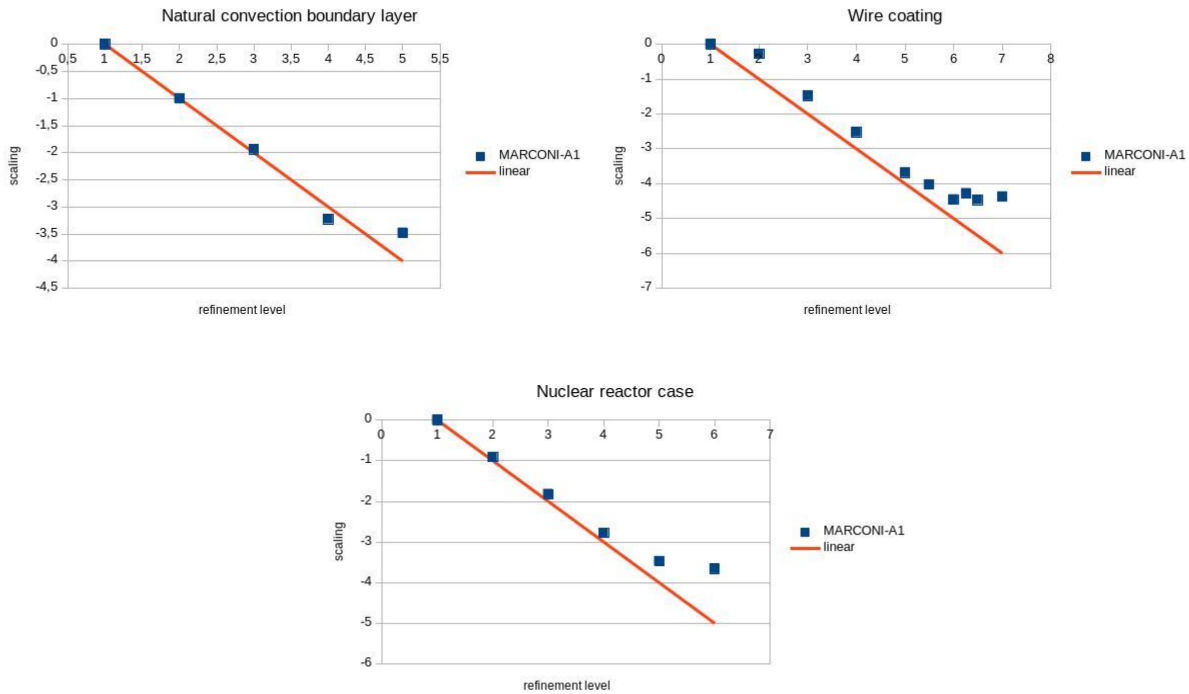

Figure 2: Effect of underlying physical modeling.

This test allowed to see that irrelevant of the size of the problem and the physical modeling considered the scaling is kept up to the approximately the same number of cells/processor. Indeed, when a processor gets less than 30 000 cells, the scaling not following any more the linear trend. While the results on Figure 2 are all performed on the CINECA cluster, the same rule of thumb was experienced on most of the architectures.

CONCLUSIONS AND LESSONS LEARNED

A detailed performance and scaling study was performed on various workstations and clusters in the Tier-0, Tier-1 and Tier-2 level. As the last part of the abstract, we would like to share some lessons we have learned while the testing:

- The I/O time is not the bottleneck in any infrastructure. It took approximately less than 5% of the overall simulation times. Therefore, except for a specific test performed on the AMD Epyc hardware (not shown here, but will be included in the presentation), we always included it in the performance tests.
- Performing re-numbering after decomposition led to up to 30% speed-up in all the tested systems. The gain, however, decayed with the increase of the number of processors used.

- We have tested on the AMD processors an OpenFOAM version that was compiled by Intel compiler and one that was with AMD compiler (AOCC 1.2). In most cases using the AMD compiler resulted in a performance gain of up to 15%.
- We could not obtain a full scaling with the AMD Epyc processor. We have performed many different tests, but in all cases after 32 processors – irrelevant of the case used – the scaling degraded, with or without SMT activated. The only difference was the speed of execution, indeed higher with SMT OFF.