HPC Benchmark Project: follow up

I. Spisso[1], R. Da Vià[2], R. Ponzini[3], F. Magugliani[4], N. Ashton[5]

[1,2,3] SuperComputing Applications and Innovation (SCAI) Department, CINECA, Italy; i.spisso@cineca.it, r.davia@cineca.it, r.ponzini@cineca.it
[4] E4 Computer Engineering SpA, Scandiano (RE) Italy, fabrizio.magugliani@e4company.com
[5]Principal CFD Specialist SA, HPC, Amazon Web Services, UK neashton@amazon.com

Table of Contents

Abstract **1**

Code Repository **1**

List of test-cases 2

Set-up of linear algebra solver: fixedITER 3

Hardware Comparison **4**

References **5**

Abstract

The current work is a follow up of the on-going activity aiming to create a framework for enabling the consistent and coherent comparison of performances using OpenFOAM technology across different HW platforms and configurations, not only in terms of pure Time to Solution (TtS) but also in terms of computational efficiency (with different KPI, i.e. Key Performance Indicators) and Energy To Solution (EtS), the latter a parameter gaining a growing importance in view of the transition to exascale.
This work shows the steps taken for creating such a framework, detailing the rationale for selecting a number of reference datasets, each exercising a specific feature of the HW platform and cluster infrastructure, as well as future platforms to enable the coherent comparison of the performance across current and new generations of HW and SW components, and the realisation of a shared repository for data and information to enable the open access to the data.

Code Repository

The code repository for the HPC TC [1] is a shared repository with relevant data-sets and information created to:

- Provide user guides and initial scripts to set-up and run different data-sets on different HPC architectures
- Provide to the community a homogeneous term of reference to compare different hardware architectures, software environments, configurations, etc.
- Define a common set of metrics/KPI to measure performances and efficiency

The benefits from an end-user perspective are:

- Opportunity for having access to a homogeneous set of data with respect to performance
- Opportunity for having access to the performance report for new HW & SW infrastructures and functionalities
- Opportunity to influence the directions of the development of the package
- Forum, information sharing

List of test-cases

The following Table 1 shows the initial list of test-cases with the specifications and the relation with HW/SW infrastructure (on-going work)

data-set	Problem size(s)	Physical features	Relation with HW/SW infrastructure				Computational density (# of cores per node)	KPIs	bottleneck(s)
			Cpu intensive	Memory intensive	GPU intensive	I/O intensive			
3-D Lid Driven Cavity flow	S, M, XL, XXL	Incompr. laminar flow, regular and uniform grid	yes	yes	no	no	Full	Time to solution Energy to solution (?) Memory Bandwidth Bound	Linear solver Algebra Data structure
HPC Motorbike	S, M, XL	…	..	..	..	..	..	..	..
ExaFOAM test-cases									

Table 1: Initial list of benchmark test-cases

- 3-D Lid Driven Cavity flow

  The first test-case chosen is the 3-D version of the Lid-driven cavity flow tutorial. This test-case has a simple geometry and boundary conditions, involving transient, isothermal, incompressible laminar flow in a three-dimensional box domain. The *icoFoam* solver is used in such test-case. It is intended to stress test the linear algebra solver, most of the time being spent in the pressure equation. In this simple geometry, all the boundaries of the box are walls. The top wall moves in the x-direction at the speed of 1 m/s while the other five are stationary.

  Three different sizes have been selected, as shown in Table 2. A further size XXL with 216 M of cells is going to be added. Figure 1 shows the computational mesh and the streamlines of U for the M test-case.

Parameters / Test-case	**S**	**M**	**XL**
Cube side length d [m]	1E-1	1E-1	1E-1
Δx [m]	1E-3	5E-4	2.5E-5
N. of cells, total (millions)	1	8	64
N. of cells, on cube edge	100	200	400
kinematic viscosity ν [m^2/s]	1E-1	1E-1	1E-1
Courant number Co	1	0.5	0.25
Reynolds number Re	10	10	10
Top wall velocity U [m/s]	1	1	1
N. of iterations	500	2000	8000
End time [s]	0.5	0.5	0.5 (0.00625)
Δt [s]	1E-3	2.5E-4	6.25E-5

Table 2: Geometrical and physical parameters and details of test- case 3D lid-driven cavity flow, size S, M and XL

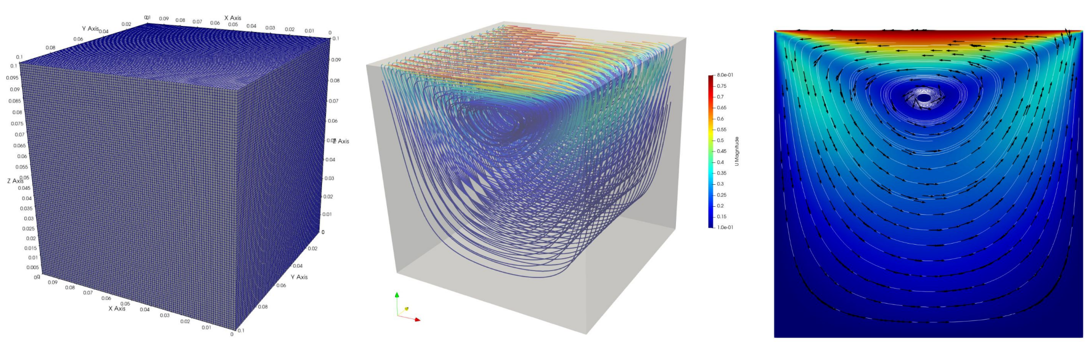

Fig. 1: Mesh and streamlines of U for the M test-case

- HPC motorbike

This test-case has been developed as an HPC extension of the standard OpenFOAM motorbike tutorial for external aerodynamics. The standard motorbike tutorial cannot be considered as an HPC benchmark since the computed mesh consists of 300k cells approximately. Three different mesh sizes are here proposed, as an HPC extension, namely S, M and XL, obtained as follows. We keep the same dimensions of the domain, i.e. a box with length L=20m, width W=8m and height H=8m, as shown in Fig.2, and increase the refinement levels with respect to the standard test. For the S mesh size we consider one half of the domain, i.e. 0<y<4, with the symmetry plane acting as a boundary side. For the M mesh size case we consider the full geometry of the wind tunnel, as the one shown in Fig. 2, while for the XL case we consider a domain that is twice the size of the M case, with two motorbikes in it. The domains of S, M and XL motorbike cases are shown in Fig. 3. We followed this approach since it allows us to increase the mesh size without refining the mesh cells. In this way the flow resolution is the same between the different cases, so the same setup, in terms of boundary conditions for velocity, pressure and turbulence fields can be used with the three different cases. As it regards the amount of mesh cells, the S, M and XL cases have 5.5, 11 and 22 million cells respectively. It is clear that bigger meshes can be computed by following the same approach, for example by putting three domains one next to the other, obtaining a 44 million cell geometry.

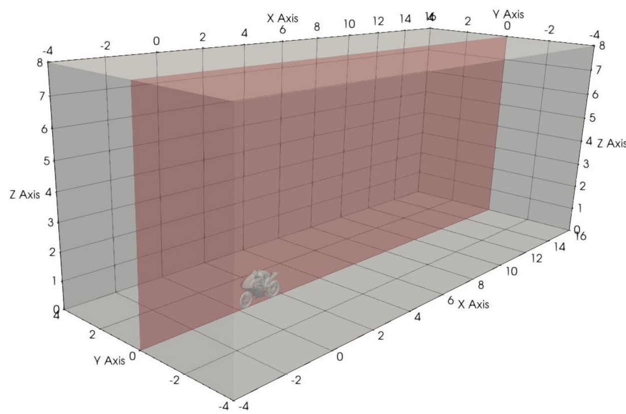

Fig. 2: reference domain for the motorbike test, with highlighted symmetry plane at y=0.

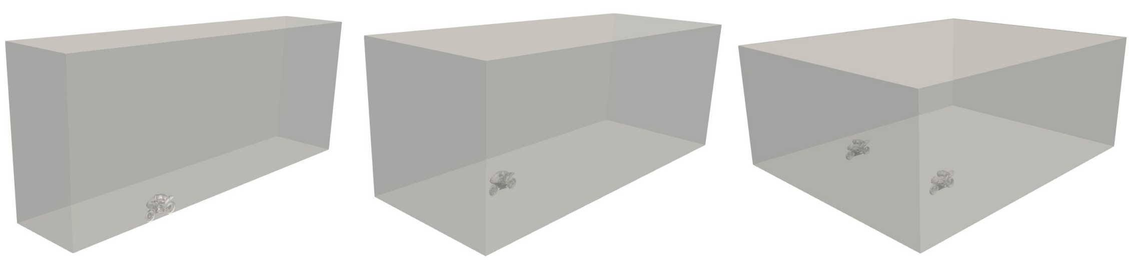

Fig. 3: from left to right, the computational domains for the S, M and XL HPC motorbike cases, respectively..

This approach is suitable for testing both strong and weak scaling of the OpenFOAM solvers.

- Industrial test-cases from WP2 of exaFOAM

During the forthcoming project exaFOAM "Exploitation of Exascale Systems for Open-Source Computational Fluid Dynamics by Mainstream Industry" selected for funding by the EuroHPC JU, many industrial stakeholders of the project have committed themselves to provide industrial HPC benchmarks to the project. The applications and grand challenges data sets will be derived from industry-based examples specified by the Observer Partners, representing the following application sectors: aerospace, power generation, automotive, environmental (volcanology), disaster (fire), polymer processing. After a careful Verification and Validation process, this benchmark will be included in the HPC repo.

Set-up of linear algebra solver: *fixedITER*

As reported in [1], different set-ups for the linear algebra solvers have been selected. For HPC benchmark and hardware comparison, it is important to select the *fixedITER* convergence criteria as shown in the set-up configuration of Fig. 4. The set-up has been chosen to have a constant computational load at each time step; the set-up has been carefully chosen so that the maximum number of iterations is reached at each time step both for pressure and the velocity's components. In this case, the computational load is fixed when running with the given solver. This case is useful for comparing different hardware configurations by keeping constant the computational load, as shown in Fig. 4 for the residuals and the number of iteration for pressure and velocity's components.

system/fvSolution

Test-case	S	M	XL
pressure			
solver	PCG	PCG	PCG
tolerance	0.00E+00	0	0
relTol	0	0	0
maxIter	250	650	1300
num. iter Tot	125,000	1,300,000	10,400,000
pFinal	-	-	
relTol	0	0	0
U velocity			
solver	PBiCGStab	PBiCGStab	PBiCGStab
preconditioner	DILU	DILU	DILU
tolerance	0	0	0
relTol	0	0	0
maxIter	5	5	5
num. iter Tot	2,500	10,000	40,000

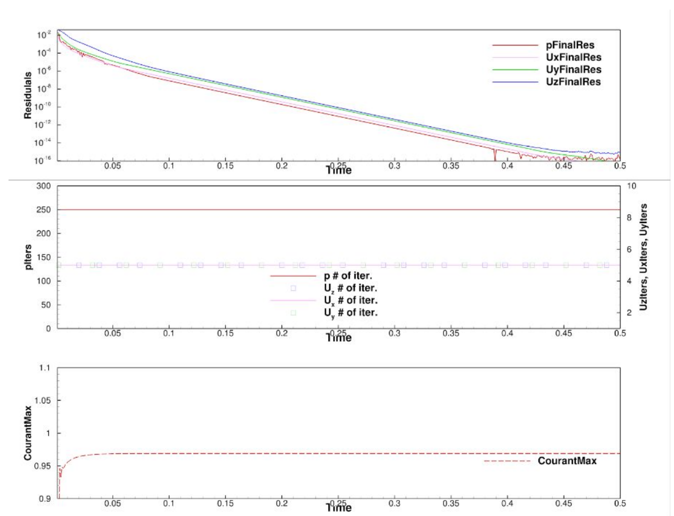

Figure 4: Set-up of solvers and time evolution of the residuals and the number of iteration for pressure and velocity's components.

Hardware Comparison

Table 3 shows the example of reporting a system configuration and the associated results.

System configuration		
Specify bare metal / Virtual Machine		
HW	CPU (model, # of cores and speed)	
	Number of CPUs per node	
	Number of nodes	
	Number of OpenFOAM processes per node	
	Memory and type per node	
SW	Complier used	

Case	Execution Time (last time)								
Computational set-up									
XL # 1									
XL #2									
XXL #1									

Table 3: Example of the report of the benchmarking system for hardware comparison

Additional files requested are:

- log of runs
- plot/table with residual (work in progress for post-process scripts)

Figure 5 shows the comparison of two different hardware architectures by using the same number of cores. The runs have been executed using the full computational density for the corresponding architecture, in this case 36 cores per Intel Xeon E5-2697 a.ka. Broadwell (Galielo's cluster), versus 64 cores for MARVELL TX2.

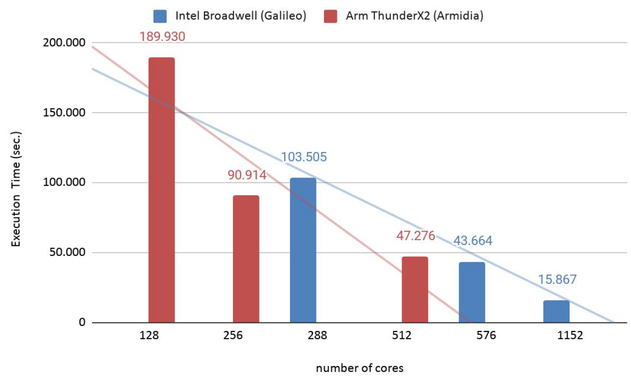

Figure 5: HPC comparison between Armidia (TX2) and Galileo (Broadwell) cluster, XL test-case. Continuous lines represent the trend line.

References

1. https://develop.openfoam.com/committees/hpc/-/tree/master