Quantum Espresso performance on ENEA and JSC HPC infrastructures.

Simone Giusepponi[1]*, Michele Gusso[2], Massimo Celino[1],
Urs Aeberhard[3] and Philippe Czaja[3]

[1] *ENEA, C. R. Casaccia, via Anguillarese 301, 00123 Rome, Italy*

[2] *ENEA, C. R. Brindisi, S.S. 7 Appia km 706, 72100 Brindisi, Italy*

[3] *IEK-5 Photovoltaik, Forschungszentrum Jülich, D-52425 Jülich, Germany*

Abstract. In this paper we report results from a benchmark in which we compare the performances of three HPC clusters: CRESCO4 and CRESCO5 located in the Portci ENEA centre, and JURECA located in Jülich Supercomputing Centre.

1 Introduction

In view of a intensive use of computational resources for an ab-initio study of a crystalline Silicon (c-Si) and hydrogenated amorphous Silicon (a-Si:H) interface (see Fig. 1), a measure of the performances of three HPC infrastructures (CRESCO4, CRESCO5, JURECA) was carried out. This enable us to set the optimal computational resources and give an estimation of the computational time. The benchmark is based on the time required to execute a self consistent calculation, for a single energy minimization in the computation of the ground-state energy of the c-Si/a-Si:H interface depicted in Fig. 1. The calculation is performed using the ab-initio PWscf (Plane-Wave Self-Consistent Field) code of the Quantum ESPRESSO suite [1, 2].

The ENEA CRESCO computing facilities are based on the multi-core x86_64 architecture and is made up of various clusters: we tested the CRESCO4 and CRESCO5 clusters located in Portci ENEA centre [3, 4].

CRESCO4 is composed by 304 compute nodes. Each node hosts a dual-socket 8 cores CPU Intel E5-2670 processor (Sandy Bridge architecture) which operates at a clock frequency of 2.6 GHz and has 20 MB of cache and avx instructions set. The total number of cores is 4864.

CRESCO5 is composed by 40 compute nodes. Each node hosts a dual-socket 8 cores CPU Intel E5-2630 v3 processor (Haswell architecture) which operates at a clock frequency of 2.4 GHz and has 20 MB of cache and avx 2.0 instructions set. The total number of cores is 640.

Both clusters have 64 GB of RAM memory per node (4 GB per core) and are interconnected via an IB 4xQDR Qlogic/Intel 12800-180 switch. The Operating System for the clusters is the Red Hat Enterprise Linux ver. 2.6.32-220.7.1.el6.x86_64. Quantum Espresso ver. 5.1.2 is compiled with intel compiler ver. 14.0 using the mathematical library MKL ver. 11.1 and OpenMPI ver. 1.4.3.

The third cluster is JURECA HPC hosted at Jülich Supercomputing Centre (JSC) [5].

*Corresponding author. E-mail: simone.giusepponi@enea.it.

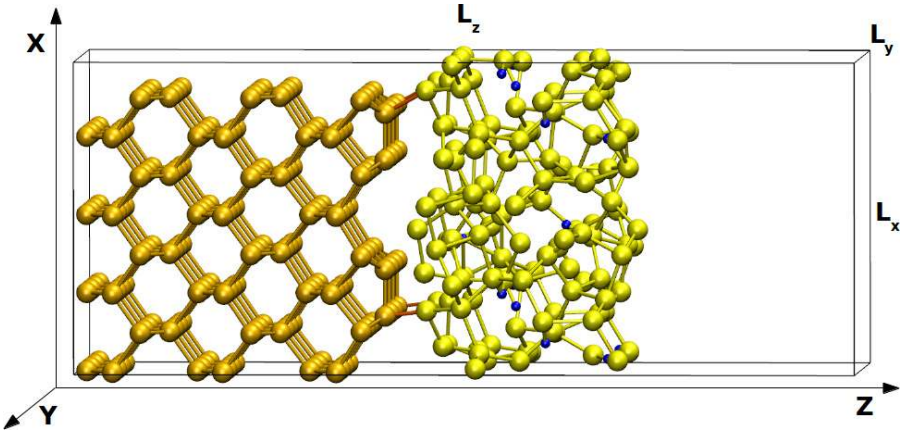

Figure 1: Snapshot of the initial a-Si:H/c-Si interface in the simulation box. The structure is infinitely extended in both x and y directions. A void region is considered to suppress the interaction, due to periodic boundary conditions, between the external surfaces. Free surfaces and a-Si:H/c-Si interface are perpendicular to the y axis. Hydrogen atoms are in blue, Silicon atoms are in orange in the c-Si side and are in yellow in the a-Si:H side.

JURECA is composed by 1872 compute nodes. Each node hosts a dual-socket 12 cores CPU Intel E5-2680 v3 processor (Haswell architecture) which operates at a clock frequency of 2.5 GHz and has 30 MB of cache and avx 2.0 instructions set. The total number of cores is 44928, the cluster has 128 GB of RAM memory per node that are interconnected via a Mellanox EDR InfiniBand. The Operating System for the cluster is the CentOS 7 Linux ver. 3.10.0-327.10.1.el7.x86_64. Quantum Espresso ver. 5.1.1 is compiled with intel compiler ver. 15.0 using the mathematical library MKL ver. 11.2 and ParaStationMPI ver. 5.0.

2 Results and discussion

To compare the performance and scalability of the three clusters, we collected the time necessary to perform a self consistent calculation, for a single energy minimization in the computation of the ground-state energy of the c-Si/a-Si:H interface depicted in Fig. 1. Total cpu and wall times for the three clusters using 48, 96, 192 and 384 cores are reported in Tab. 1 and illustrated in panel a) of Fig. 2. The calculation is performed considering only the Γ point and using MPI parallelization.

Moreover, in Fig. 2, there are shown the speed up (panel b) and the resulting efficiency (panel c) calculated for each cluster with respect to the time measured in the 48 cores execution. This type of camputation has a good parallel scalability until 192 cores with an efficiency of about 75%.

In addition, in Tab. 1 for each row, it is reported the percent gain in time (numbers in parenthesis) for CRESCO5 and JURECA with respect to the CRESCO4 execution time. The results confirm that the more recent processors, E5-2630 v3 and E5-2680 v3 mounted on CRESCO5 and JURECA clusters, respectively, have the best performance with reductions of execution times, in the range 30÷40% for JURECA and in the range 10÷20% for CRESCO5. Then, even if, they have the same architecture (Haswell) and the same instruction set (avx 2.0), E5-2680 v3, being two and half times more expensive, legitimises its cost having better performance. This is also emphasised in panels b) and c) of Fig. 3, in which the speed up and efficiency were calculated with respect to the execution time for 48 cores on CRESCO4 cluster.

Finally, in the last two colums of Tab. 1, the difference between wall and cpu times, and the corresponding percent variations are shown. Data indicate that CRESCO4 and CRESCO5 have differences in time less

Table 1: Total cpu and wall time to perform a single energy minimization in the computation of the ground-state energy for the c-Si/a-Si:H interface, at different number of cores used in the calculation. Differences between wall and cpu times and the corresponding percent variations are shown in the last two columns. Moreover, for each row, the percent gain of times on CRESCO5 and JURECA clusters with respect to the time on CRESCO4 cluster are reported.

cores	cluster	Total cpu time t_1(s)	Total wall time t_2(s)	$t_1 - t_2$(s)	$\frac{t_2 - t_1}{t_1}$(%)
48	CRESCO4	966.3	1037.8	71.5	7.4
	CRESCO5	889.2 (-8.0%)	957.5 (-7.7%)	68.3	7.7
	JURECA	695.7 (-28%)	743.1 (-28%)	47.4	6.8
96	CRESCO4	571.7	610.4	38.7	6.8
	CRESCO5	483.1 (-15%)	520.1 (-15%)	37.0	7.7
	JURECA	346.8 (-39%)	411.7 (-33%)	64.9	19
192	CRESCO4	371.7	389.9	18.2	4.9
	CRESCO5	306.5 (-18%)	317.9 (-18%)	11.4	3.7
	JURECA	211.5 (-43%)	277.6 (-29%)	66.1	31
384	CRESCO4	291.3	306.7	15.4	5.3
	CRESCO5	224.9 (-23%)	238.2 (-22%)	13.3	5.9
	JURECA	180.9 (-38%)	253.2 (-17%)	72.3	40

than 8%, in fact, they use the same interconnection network. On the contrary, JURECA has differences in time that enlarge with the number of cores used in the calculation: from about 7% with 48 cores up to 40% at 384 cores. This is maybe due to the size of the cluster (1872 compute nodes and 44928 cores), in which there are hundreds/thousands concurrent jobs.

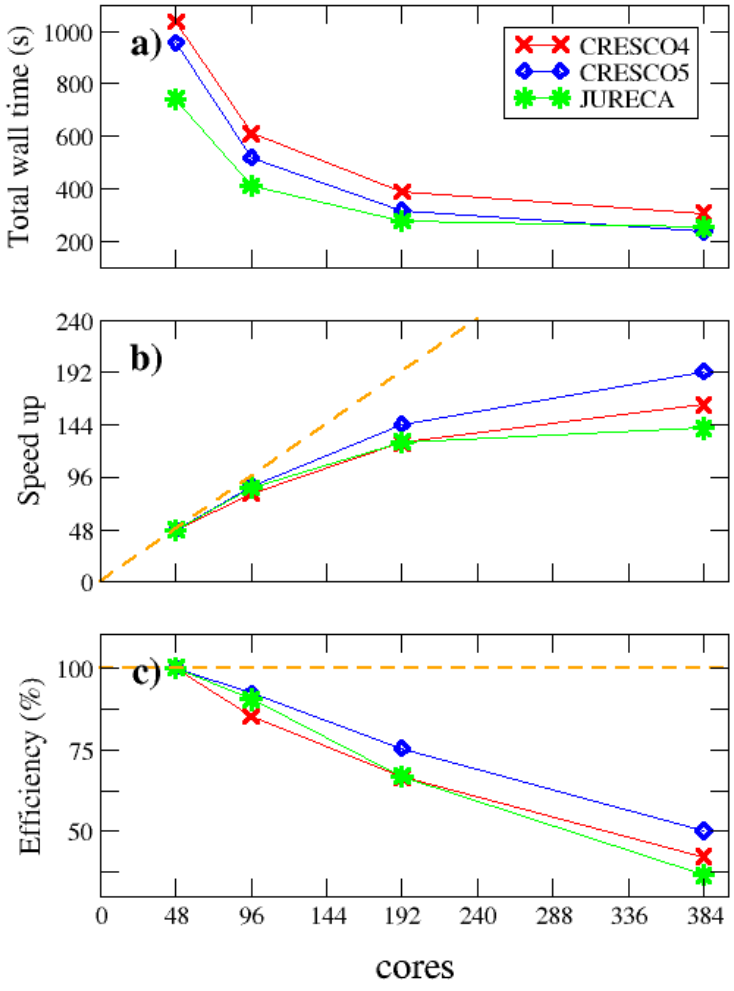

Figure 2: a) Total wall time (s) for a self consistent calculations for a single energy minimization. b) Speed up calculated with respect to the time in the 48 cores calculation, and c) the corresponding efficiency (%). Red lines(symbols) are for CRESCO4, blue lines(symbols) are for CRESCO5 and green lines(symbols) are for JURECA. Dashed orange lines represent ideal parallel scalability.

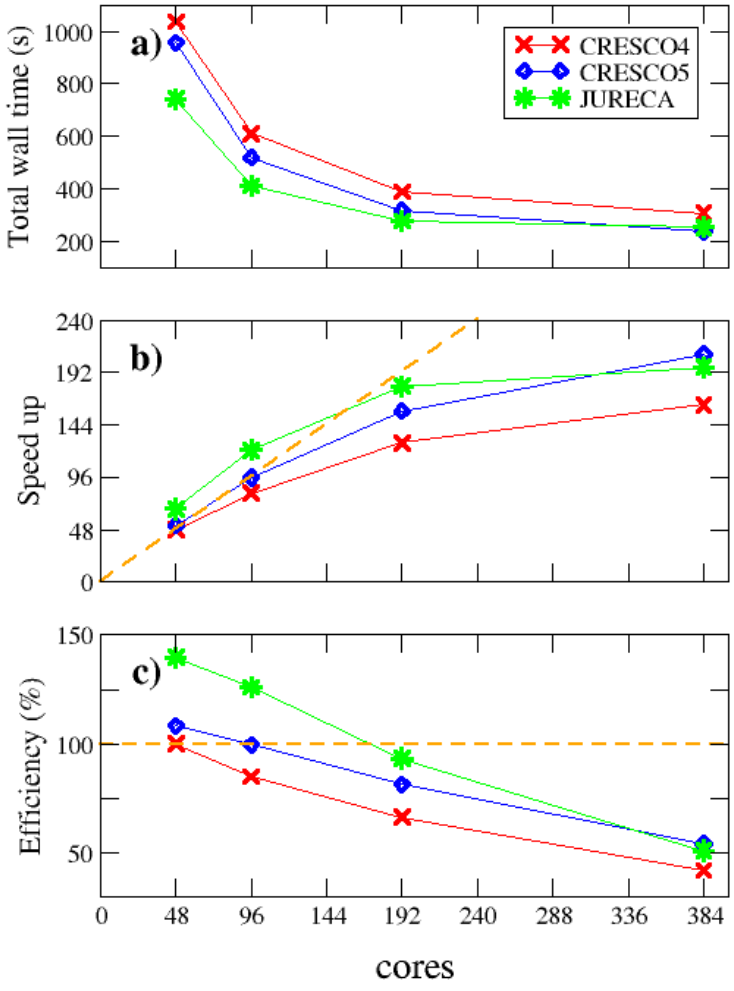

Figure 3: a) Total wall time (s) for a self consistent calculations for a single energy minimization. b) Speed up calculated with respect to the time in the CRESCO4 48 cores calculation, and c) the corresponding efficiency (%). Red lines(symbols) are for CRESCO4, blue lines(symbols) are for CRESCO5 and green lines(symbols) are for JURECA. Dashed orange lines represent ideal parallel scalability.

References

[1] P. Giannozzi *et al. J. Phys.: Condens. Matter.*, 21:395502, 2009.

[2] www.quantum-espresso.org.

[3] www.cresco.enea.it.

[4] www.eneagrid.enea.it.

[5] www.fz-juelich.de/ias/jsc/en/expertise/supercomputers/jureca/jureca_node.html.