Porting basic linear algebra algorithms on the EGI parallel platform

A. Costantini[1,2], O. Gervasi[2], L. Pacifici[1], S. Crocchianti[1], C. Manuali[1], A. Laganà[1],

[1] Department of Chemistry, University of Perugia (Italy)
[2] Department of Mathematics and Computer Science, University of Perugia (Italy)

1) Porting of linear algebra algorithmsc

The increasing availability of computer power on distributed platforms makes it easier to implement basic linear algebra algorithms that are level complexity codes taken from typical library routines or algorithms popular among the Molecular and Material Science (MMS) community members.
We outline the work carried out in our laboratory by porting on the EGI Infrastructure a set of these algorithms and performing MPI calculations to test the parallel performance of the used platform. cvcv

2) EGI Grid platform - Test case

All the Grid nodes that support COMPCHEM VO have been tested using three different matrix multiplication algorithms

- (Cannon algorithm, Fox algorithm and Strassen algorithm)

From a preliminary analysis performed by running the related JDL file in which the requirements MPI-START & MPICH have been specified, it was found that 16 over a total of 25 sites supporting COMPCHEM VO support also MPI applications.

The performance of each site has been obtained running the codes sequentially in one node and in parallel in 2 ,4 ,8, 16, 32, 64 nodes on the same cluster
The global performances and the statistical analysis carried out by submitting MPI jobs have been compared with those obtained in 2009.

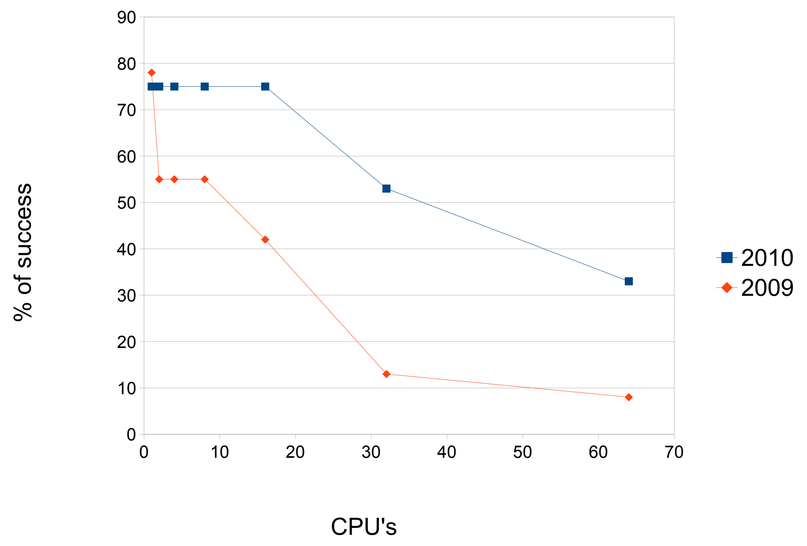

The reduction in sites supporting MPI registered in EGI (in 2009 were 22/25 in EGEE) is basically due to the introduction of the MPI-SAM tests (now NAGIOS tests). The tests now assure the basic requirements for a job submitted with MPI tags.

3) EGI Grid platform - Statistics

From the tests performed by running the parallel MPI version of a set of common matrix multiplication algorithms in the EGI environment has emerged that the number of jobs requiring up to 8 CPUs and terminated successfully has increased of about 20% and the number of jobs requiring from 16 up to 64 CPUs has increased of about 30% if compared with 2009 results. The obtained results indicate an increase of reliability on the use of MPI on the present Grid environment.

Up to 8 CPUs

Job status	%(2009)	%(2010)
Success	53	75
Not success	47	25

Not success	%(2009)	%(2010)
Abortion (CE)	52	0
Scheduler Error	39	100
MPI-START	9	0

From 16 to 64 CPUs

Job status	%(2009)	%(2010)
Success	21	54
Not success	79	46

Not success	%(2009)	%(2010)
Abortion (CE)	73	0
Scheduler Error	23	93
Proxy expired	4	7

5) Performances of the evaluated platform

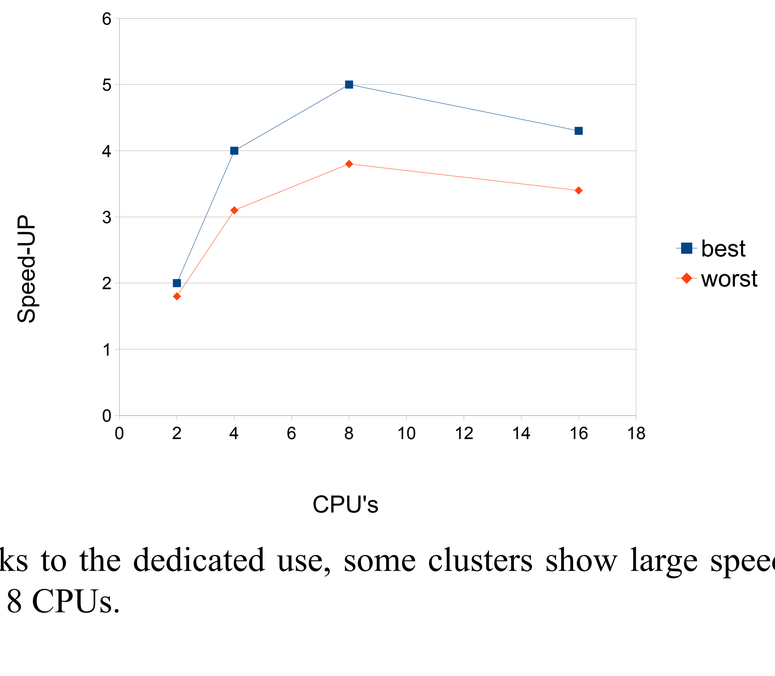

Thanks to the dedicated use, some clusters show large speed-ups up to 8 CPUs.

Conclusions

The porting of basic linear algebra algorithms onto the Grid infrastructures can be seen as part of a more general effort to build a solid platform to implement the parallel versions of complex suites of codes. The presented case study demonstrates the possibility of using the parallel capabilities of the European Grid and can be used as an example for those Communities which are interested in the porting of their parallel applications into the Grid environment.