

Main Message:

In June 2011 the MPI subtask members produced input for the EGI-InSPIRE PY1 review¹.

As per the CCMST/UNIPG workplan, UNIPG produced a report entitled “WORK PERFORMED BY THE RESEARCH CLUSTER OF THE UNIVERSITY OF PERUGIA: MPI on the Grid, Second Trimester Report” in July 2011. This covers work on using a hybrid Cloud/Grid model to enable GPU resource allocation and exploitation. UNIPG has also continued work on the following:

- parallelisation of some linear algebra routines,
- implementation of the “horse race” programs for atom diatom quantum reactive scattering of UNIPG ABC and RWAVEPR,
- Evaluation of Chimere in preparation for parallelisation.
- Parallelisation of DL_POLY, NAMD and VENU96.

TCD has installed two new sub-clusters to support Parallel Jobs. The first is a cluster of gLite 3.2 enabled Sony Playstations. This exploits the “Cell Broadband Engine” processor. For some codes this offers approximately a twenty times speedup over conventional processors. The second sub-cluster is a cluster of 32 GP-GPUs. It consists of 16 nodes, each with 2 NVIDIA GPUs. This cluster uses CUDA 4.0 libraries.

CSIC/TCD evaluated early releases of gLite-MPI and WMS-3.3. The products are not yet ready for production.

Issues and Mitigation:

Support for the prevalent batch system “Torque/Maui” on the EGI infrastructure increases the chances of delays to the timely production of updated MPI related RPMs.

¹ <https://www.egi.eu/indico/getFile.py/access?contribId=8&resId=0&materialId=slides&confId=475>

Support for GPU resources is currently not available in the MAUI scheduler. There is limited support for GPU resources in Torque 2.5 onwards, and SGE has some scheduling mechanisms for GPUs. There are issues in granting exclusive access to the GPU resource on multi-core systems. A possible solution is to use virtualization. This ensures only one job has access to the GPU resource. This has been tested by UNIPG², and TCD will use the StatusLab cloud distribution software for GPU provisioning.

Delivery of UMD gLite-WMS 3.3 is late. This should have support for user-defined allocation of processes per node, and will allow for better allocation of resources for MPI and other parallel jobs.

Conferences and Papers:

Enol Fernandez, "Support to MPI applications on the Grid" at 5th Iberian Grid Infrastructure Conference (IBERGRID 2011), Santander, June 2011.

Antonio Lagana, presented at ICCSA 2011, Santander, Spain. June 2011.

For publications, see footnote 2.

² See R. Baraglia, M. Bravi, G. Capannini, A. Lagana', E. Zambonini, A Parallel code for time independent quantum reactive scattering on CPU-GPU platform, Lecture Notes Computer Science 6784, 412-427 (2011); and

L. Pacifici, D. Nalli, D. Skouteris, A. Lagana', Time dependent quantum reactive scattering on GPU, Lecture Notes Computer Science 6784, 428-441 (2011)

Plans for next quarter:

- SA3 “Software Roadmap” input
- EGI MPI user and site administrator surveys,
- UNICORE and ARC MPI documentation,
- Continued work on the MPI cookbook,
- MPI workbenches for Computational Chemistry and Fusion Communities:
 - GROMACS, CHIMERE, GAMESS
- Preparation for UMD gLite-MPI and glite-WMS 3.3 releases.