A test bed for developing a CMMST grid synergistic model

Leonardo Pacifici[1], Carlo Manuali[1], Alessandro Costantini[2], Giuseppe Vitillaro[3], Antonio Laganà[1],
1. Department of Chemistry, Biology and Biotechnologies, University of Perugia
2. IGI-INFN, Perugia, Italy
3 ISTM-CNR, Perugia

1.   Introduction

The Chemistry, Molecular, Materials Sciences & Technologies (CMMST) is a Virtual Research Community (VRC) established thanks to the activities of the EGI-InSPIRE Virtual Team (VT) project "Towards a CMMST VRC" (see ref [1]) promoted by the COMPCHEM Virtual Organization (VO) in a joint endeavour with other VOs, computational Molecular Science research projects and the participation of some Associations and spinoffs.

An important aim of the CMMST VRC is the development and adoption of a truly collaborative/competitive (synergistic) distributed computing model in which not only the barriers between compute platforms are lowered but also the difference between users and producers is weakened, programs designed and implemented for different research areas interoperate and collaboration is encouraged and rewarded in terms of its quality parameters to the end of fostering sustainability through a higher level of competitiveness. This project is a great leap forward from the opportunistic model (that allows the user to utilize the resources made available by the providers) to a model that prompts a proactive effort of the users to develop tools and environments rewarding such attitude. In the present paper we sketch the steps undertaken in this direction by a group of users of Perugia on behalf of the local community and, more in general, of the CMMST VRC.

2.   The sharing of expertise

The first practical step undertaken by the CMMST community in Perugia has been the establishing of a cluster of competences and expertise in Computational Molecular and Materials Sciences and Technologies. Such cluster, called Scientific Computing Cluster of the University of Perugia (SCC-UP), gathers together autonomous research teams (which remain autonomous) of the University of Perugia (from the Chemistry, Biology and Biotechnologies and the Mathematics and Computer Science Departments, the Perugia INSTM-CNR section and the Perugia INFN section. The infrastructure is meant to be open to the participation of other Universities, research Institutions or private sector investigators of the field.

As is at its present composition the cluster relies on the expertise on:

- theoretical and computational modelling of atomic and molecular processes both at quantum and classical mechanics level as well as distributed computing of the Computational Dynamics and Kinetics (CDK) group of the University of Perugia;

- theoretical and computational modelling in various fields of inorganic chemistry applications of the Theoretical and Computational Inorganic Chemistry (TCIC) group of the University of Perugia;

- computer and computational sciences in knowledge management systems handling in networked applications of the High Performance Computing (HPC) group of the University of Perugia;

- Materials and nanostructure computational design for modern applications (including photo-processes, catalysis, cultural heritage, etc.) group of the Perugia section of the Istituto di Scienze e Tecnologie dei Materiali (ISTM);

- grid management and training and user support of the Italian Grid Infrastructure (IGI) sector operated by the Istituto Nazionale di Fisica Nucleare (INFN) in charge to the Perugia section;
The various groups have developed (Bertha, ABC, RWAVEPR) and/or manage locally (GAMESS-US, NWCHEM, QUANTM ESPRESSO, DL_POLY) various programs and packages (mainly Open Source) specialized in Sciences and Technologies of Molecular and Materials Sciences including the service oriented Grid Empowered Molecular Simulator (GEMS) as well as the Grid Learning Object REPository (GLOREP)

3.   Hardware

The second practical step undertaken by the CMMST community in Perugia has been the assemblage of a suitable test bed for the implementation of such synergistic distributed computing model at hardware level made of resources owned independently by the different research teams of SCC-UP.
The machines owned by the groups are
CDK:
* 40 cores (dual core processors), Cluster of 4 workstations SUN, 18 processors, 50 TB disk space
* 200 cores (four core processors)
TCIC + ISTM
* 392 and 388 cores (from different processors with 128 Gb central memory) interconnected through independent 40Ghz switch infiniband and copper UTP gigabit ethernet.
* 2 nodes E4 7116 with Supermicro mather board with remote IPMI 2.0 management and console, 24Gb RAM DDR3-1333, dual processor Intel Xeon E 5670, 2.40Ghz, 4 core, 12Mb cache, both interfaced to an NVIDIA Tesla S2050 1U, 4GPU Fermi S2050, 3Gb RAM per GPU
* 16 core Intel Xeon 2.40Ghz (120Gflops double precision) and the development environment mvapich/openmpi for parallel message passing applications, for both peer to peer infiniband and UTP copper gigabit Ethernet links.
* 4 GPU Fermi NVIDIA S2050 (2TFlops double precision, 2GPU per node) and the CUDA development environment for applications able to exploit the high performance of the NVIDIA coprocessor.
* virtualized data centre equipped with a 3COM 3848 switch, 48 UTP ethernet copper gigabit gates, a 3COM 2924-SFP switch, 24 UTP ethernet copper gigabit gates, 2 IBM 2005-16B SAN switch, 16 16Gbits gates, 2 IBM TotalStorage DS4700 of 4.2Tb, 4Gbits SAN, 2Gbits fiber channel disks, IBM TotalStorage DS4100 of 4.6Tb, 2Gbits SAN, SATA disks, IBM TotalStorage 3573 Tape Library, IBM BladeCenter E (Type 8677), 5 IBM x346 providing the computational cluster with a network environment, a distributed storage and flexible back ups easy to adapt to different and rapidly changing needs.
HPC:
3 laboratories equipped each with 20 PCs for training use.

4.   The roadmap to future

Next steps that will be undertaken by the SCC-UP are to connect to the Grid the test bed and to implement the Grid Framework GriF [2] originally designed within the activities of the Perugia group of COMPCHEM.
GriF is a Service Oriented Architecture (SOA) Collaborative Framework designed to facilitate the use of the Distributed Compute Infrastructure (DCI) by non specialists with a particular focus on the optimization of the selection of computing elements (based on Quality of Service (QoS) parameters worked out from the information provided at run time by the grid for running single and parameter study applications. It consists on a set of java modules aimed at submitting and

2

monitoring jobs on different computing platforms and evaluate the QoS offered by the computing site and of the use made of the resources by the users.

GriF is accompanied by the introduction of a metrics aimed at evaluating the Quality of a User (QoU) and of his/her contributions of the community according to requirements set for that purpose by the community itself. Using the other tool GCreS [3] also originally designed within the activities of the Perugia group of COMPCHEM and to be further developed by the CMMST VRC, the work done by the users on behalf of the community will be rewarded by credits' assignment.

This is bound to foster a profound innovation in molecular simulations by paving the way for high performance ab initio exact treatments of both the electronic structure and nuclei dynamics of reactive processes of increasingly larger systems that are at present dealt using force field and classical mechanics (commonly referred to as molecular dynamics) techniques that ignore completely quantum effects which are of vital importance for some scientific and technological applications.

Other key problems to solve are:
The proposed simulator, though being a derivation of GEMS, is now going to be composed of software produced and/or maintained by research groups not belonging to the same virtual organization. In addition, the used software will be made of a large fraction of codes native of supercomputer environments and of middleware differing from those popular in EGI. This will require a huge effort of harmonization and standardization. Such effort will be in charge mainly to the coordinating institution that will have to identify technologies and services suitable to support the implementation of a synergistic and collaborative model among users and providers of both resources and technology. For that purpose a comprehensive support is going to be given as deliverables produced during the EGI-InSPIRE Virtual Team (VT) project "Towards a CMMST VRC". The implementation of the first prototype HPC/HTC workflow of GEMS will focus, once more, on the exact treatment of gas phase elementary reactions by making use of the WS-PGRADE technology. WS-PGRADE will be used to develop the related workflow consisting in a procedure that distributes, depending on the chosen system and parameters, either on a HTC or on a HPC platform, the generation of an appropriate set of electronic structure calculations and then launches on them a set of strongly coupled calculations to be run either on a HTC or on a HPC platform (again depending on the chosen system and parameters). The chosen prototype application will be then transformed into a HPC/HTC pre-production workflow.

Then, within a collaborative effort between SZTAKI and SCC-UP, in order to ensure access to individual grant-based usage of the resources (typically the HPC ones) or adopt an alternative community grant type, further modifications will have to be introduced into WS-PGRADE in order to allow it to deal with community access and robot certificates. WS-PGRADE accounting information will have to be modified in order to deal with collected data from the grid sensors about running time performances of the jobs and of the users in order to satisfy GriF requirements. In particular in order to formulate QoS and QoU evaluations it is necessary that the framework (grid layer) automatically records in a data base new information about users and related jobs that can be then regularly collected (say via command line) to feed the adopted metrics. Then, following an exhaustive and correct recording of such data, it is also necessary to implement a server (or a service) to evaluate the related formulae and produce the resource ranking (typically at regular intervals) for the users themselves and the management. This will be also an effort mainly in charge of SZTAKI that will extend WS-PGRADE with the resource metrics and ranking mechanisms required by the CMMST community. The above mentioned built-in broker will be able to use these information for improving its resource selection algorithms.

To the end of bridging in a synergistic fashion HPC and HTC platforms and overcome so far the highly unsatisfactory situation in which neither the use of HPC nor that of HTC alone is optimal to

3

meet the requests of a variety of complex CMMST applications (at present use is made of an experimental SSH procedure) the functionalities of the improved version of GriF will be exploited to the end of interoperating large computational applications. This in fact allows an optimization of the usage of both HTC and HPC computing resources because it is not infrequent the case in which a user utilizes HPC platforms not as such but as a bunch of loosely coupled processors underutilizing their fast dedicated network. At the same time HTC users may utilize massively distributed HPC platforms to solve tightly coupled computational tasks ending up by wasting a large amount of time in transferring data on the net. A coordination of the two types of platforms to interoperate via a single workflow (or workflow of workflows) and properly manage the various components on the most appropriate hardware, would instead allow a clever composition of complex applications optimizing the use of the various computing resources and providing the users with the best level of performance. This is, indeed, the goal of the agreement reached between EGI and XSEDE (https://www.xsede.org/) aimed at allowing to interoperate from the European Grid infrastructure on the US supercomputer network.

An important outreach of the planned effort is the chaining of the molecular simulator to some molecular data-bases and knowledge management systems of molecular structures and processes. For them appropriate representations obeying to at least de facto standards of format will be worked out. The target use of such data and knowledge is education and training purposes and the use of a synergistic model (of users and providers of both resources and technology) allows to deal with molecular data and knowledge contents by following the approach adopted for selection and quality evaluation of compute and storage elements.

One of the activities of this task will be the generation and implementation of the requirements for dealing with data and knowledge. In particular the proper metadata will be identified in order to ensure both the interoperability of the data among different packages of the simulator (especially between the modules of electronic structure and dynamics) and the support to the development and retrieval of the information for comparison of the quality of sources of different quality. This will pave the way to a knowledge exploitation by targeting both education (of senior and younger generations) and publication of the materials (including their prompt re-usability) produced by the different members of the community. This will foster the activation of a virtuous circle of continuously improved versions of relevant materials and their efficient use for training and learning. This will also foster the development of editorial initiatives facilitating the re-use of published materials.

Finally, the CMMST-CC will target the promotion at the same time of molecular sciences and of distributed ICT technologies competences dealt in a synergistic fashion by the community members. In this task, the CMMST-CC will aim at disseminating the project activities and establish sustainability measures. This will focus in particular on extending competences and services to other communities and linking with SMEs for science transfer into innovative technologies.

These measures will target in a diffuse way not only all the members of the CMMST VRC but also related projects and Associations. They will also target some SMEs operating in the field of ICT based education and training technologies recently networked under the umbrella of the Virtual Education Community standing committee of the ECTN Association. The goal is the development of an operational modality aimed at ensuring sustainability to the Consortium beyond the lifetime of the present project. Particular effort will be paid for establishing links with spinoffs and SMEs operating in science transfer into innovative technologies. To this end specific channels facilitating the export of services to external communities and enterprises (with particular relevance for environmental agencies, research based education institutions, pharmaceutical companies, molecular biology and biotechnologies, etc.) will be investigated for activation. The strategy that will be elaborated will focus on providing a financial basis for establishing links with some start-

ups and spin-offs in particular by 1) promoting a Consortium with chemistry and ICT based companies and institutions; 2) establishing a user panel aimed at advising the Consortium on the recent requirements and interests of the society; 3) developing a sustainability model.

6. Conclusions

In the present document the motivations and the guiding ideas for establishing a local SCC-UP cluster are discussed. The document depicts the present scenario of the local computational resources and of their future usage by analysing the characteristics of the existing compute time allocation model as opposed to the synergistic one to be adopted by the proposed initiative. The document enumerates also the technical and non technical aspects of such model with relation to the fostering of the selection of the resources (from personal systems to supercomputers) and services (from number crunching to massive data handling on heterogeneous platforms) as well as the advanced usage of QoS and QoU for introducing community economy based on a credit system will produce. On this ground, the SCC-UP will be able not only to share hardware and software to run out higher level of complexity application but will act to support users intending to provide services. This will have a twofold effect:

a) enhance collaboration to the end of developing real-like simulations using stable versions of their programs as components to be offered as a service to the other users;

b) encourage competition to the end of making the data produced the most accessible possible and the services provided the most efficient possible so as to exploit their use;

in a clearly sustainable way allowing also the re-use of the produced data in science as well as in knowledge handling for education, training, dissemination and the development of big data science to which chemistry is ancillary.

REFERENCES

[1] https://wiki.egi.eu/wiki/Towards_a_CMMST_VRC

[2] C. Manuali, A. Lagana', GRIF: A New Collaborative Framework for a Web Service Approach to Grid Empowered Calculations, Future Generation of Computer Systems, 27(3), 315-318 (2011).

[3] C. Manuali, A. Lagana', A Grid Credit System Empowering Virtual Research Communities Sustainability, Lecture Notes Computer Science, 6784, 397-411 (2011).

[4] E. Rossi, S. Evangelisti, A. Laganà, A. Monari, S. Rampino, M. Verdicchio, K. Baldridge, G.L. Bendazzoli, S. Borini, R. Cimiraglia, C. Angeli, P. Kallay, H.P. Lüthi, K. Ruud, J. Sanchez-Marin, A. Scemama, P. Szalay, A. Tajti, Code Interoperability and Standard Data Formats in Quantum Chemistry and Quantum Dynamics: the Q5/D5cost Data Model, J. Comp. Chem. Journal of Computational Chemistry 35(8), 611-621 (2014).