

THE INTERACTIVE EUROPEAN GRID: PROJECT OBJECTIVES AND ACHIEVEMENTS

J. Marco^a, I. Campos^a, I. Coterillo^a, I. Díaz^a, A. López^a, R. Marco^a,
C. Martínez-Rivero^a, P. Orviz^a, D. Rodríguez^a, J. Gomes^b, G. Borges^b, M. Montecelo^b,
M. David^b, B. Silva^b, N. Dias^b, J. P. Martins^b, C. Fernández^c, L. García-Tarrés^c,
C. Veiga^c, D. Cordero^c, J. López Cacheiro^c, I. López^c, J. García-Tobío^c, N. Costas^c,
J. C. Mourião^c, A. Gómez^c, W. Bogacki^k, N. Meyer^k, M. Owsiak^k, M. Płóciennik^k,
M. Pospieszny^k, M. Zawadzki^k, A. Hammad^d, M. Hardt^d, E. Fernández^e, E. Heymann^e,
M. A. Senar^e, A. Padee^f, K. Nawrocki^f, W. Wislicki^f, P. Heinzlreiter^g,
M. Baumgartner^g, H. Rosmanith^g, D. Kranzmüller^g, J. Volkert^g, S. Kenny^h,
B. Coghlan^h, R. Pajakⁱ, Z. Mosurskaⁱ, T. Szymochaⁱ, P. Lasonⁱ, L. Skitalⁱ, W. Funikaⁱ,
K. Korcylⁱ, J. Pieczykolanⁱ, K. Bałosⁱ, R. Słotaⁱ, K. Guzyⁱ, L. Dutkaⁱ, J. Kitowskiⁱ,
K. Zielińskiⁱ, L. Hluchý^j, M. Dobrucký^j, B. Šimo^j, O. Habala^j, J. Astalos^j, M. Ciglan^j,
V. Sípková^j, M. Babík^j, E. Gatiaľ^j, R. Vallés^l, J.M. Reynolds^l, F. Serrano^l,
A. Tarancón^m, J.L. Velasco^m, F. Castejónⁿ, K. Dichev^o, R. Keller^o, and S. Stork^p

- a) *Instituto de Física de Cantabria, IFCA, CSIC-UC, Santander, Spain*
 - b) *Laboratório de Instrumentação e Física Experimental de Partículas, Lisboa, Portugal*
 - c) *Centro de Supercomputación de Galicia, Santiago de Compostela, Spain*
 - d) *Steinbuch Centre for Computing, Forschungszentrum Karlsruhe, Karlsruhe, Germany*
 - e) *Universidad Autónoma de Barcelona, Barcelona, Spain*
 - f) *Interdisciplinary Centre for Mathematical and Computational Modelling, University of Warsaw, Warsaw, Poland*
 - g) *GUP Linz, Joh. Kepler University, Linz, Austria*
 - h) *Trinity College Dublin, Dublin, Ireland*
 - i) *Academic Computer Centre CYFRONET AGH, Krakow, Poland*
 - j) *Institute of Informatics, Slovak Academy of Sciences, Dúbravská cesta 9, 845 07 Bratislava, Slovakia*
 - k) *Poznan Supercomputing and Networking Center, Poznan, Poland*
 - l) *Instituto de Biocomputación y Física de Sistemas Complejos-BIFI, Zaragoza, Spain*
 - m) *Department of Theoretical Physics and BIFI, University of Zaragoza, Spain*
 - n) *Laboratorio Nacional de Fusión and BIFI, Asociación Euratom/Ciemat, Madrid, Spain*
 - o) *High-Performance Computing Center Stuttgart, Stuttgart, Germany*
 - p) *University of Coimbra, Portugal*
- e-mail: {marco, iscampos}@ifca.unican.es

Abstract. The Interactive European Grid (*i2g*) project has set up an advanced e-Infrastructure in the European Research Area specifically oriented to support the friendly execution of demanding interactive applications. While interoperable with existing large e-Infrastructures like EGEE, *i2g* software supports execution of parallel applications in interactive mode including powerful visualization and application steering. This article describes the strategy followed, the key technical achievements, examples of applications that benefit from this infrastructure and the sustainable model proposed for the future.

Keywords: Grid technologies, interactivity, visualization, parallel computing on the grid

1 INTRODUCTION

Grid computing technology[1] has made possible the setup of large scale distributed e-infrastructures like those supported by the EGEE [2] project in Europe or OSG [3] in USA. They provide a large volume of computing resources (more than 30 000 processors) for distributed communities of researchers collaborating in international projects, enabling an e-Science activity.

The basic objective of the EGEE project, and also of the precedent EDG [4] project, is to support the distributed processing of very large data volumes, like the experimental data that will be recorded at the Large Hadron Collider at CERN. This processing, and the execution of corresponding simulations, are executed in batch mode. The EGEE project developed the gLite [5] grid middleware as an extension to existing basic packages such as Globus [6] to satisfy these requirements. The gLite distribution is widely used now over the EGEE e-Infrastructure and others related to it in Europe.

However, the EGEE project does not address specifically an important point in the e-Science needs of researchers: to establish the final results of many studies, they need to perform complex interactive analysis and simulations that may include visualization, parameters tuning, etc.

As described in this paper, many examples can be found in different research areas with some common factors:

- researchers would like to see the results in minutes rather than in hours, even if they have to process a significant data volume or apply complex algorithms. They expect that the e-Infrastructure will provide an added value at the “tera” level (teraflop-terabyte)
- the users need a friendly but flexible and powerful interface, similar to their desktop environment.

The Interactive European Grid project started in May 2006 and extends over two years, in the context of the EU 6th Research Framework Program. The project scope is to deploy a production quality e-Infrastructure, interoperable with those using glite, but providing advanced support for the scientific applications, in particular interactivity, parallel execution and graphical visualization.

The roots of the project date back to the CrossGrid [7] project (2002–2005), where a first job scheduler (the Crossbroker) was developed to manage both interactive and batch jobs of both sequential and parallel types. From the final user point of view the project introduced a Java Webstart application, the Migrating Desktop [8] that enables a user-friendly access from any platform.

The Interactive European Grid project has exploited and consolidated these achievements, to provide a production infrastructure interoperable with glite based resources supporting advanced application execution.

The project activities and their dependencies are shown in Figure 1.

An application support activity is dedicated to serve as a bridge between the middleware developers and the user communities: the team has analyzed the requirements of the reference project applications and gave feedback to infrastructure and middleware developers.

The middleware activity is in charge of developing and adapting existing middleware to support interactivity, parallel job execution and visualization. In this respect a substantial effort goes into keeping interoperability with glite infrastructures.

Finally there is the setup and operation of the production quality grid infrastructure where the *i2g* middleware is deployed, to provide the adequate service as required by researchers.

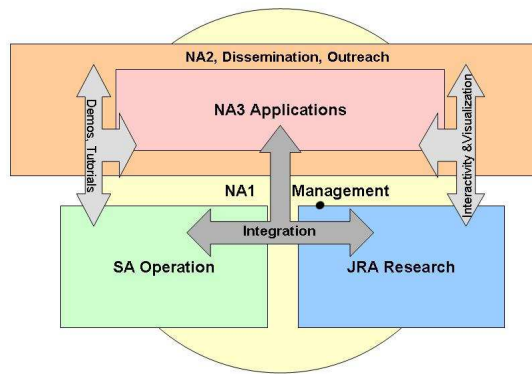


Fig. 1. Project activities definition and intersections

Operating a production infrastructure implies a high level of reliability so that researchers get convinced that the Grid is a mature technology which they can trust and rely on for the everyday work. The challenge concerns both middleware and tools developers, and also the e-infrastructure deployment teams. The close collaboration between the different activities involved in the project has been a fundamental aspect. Middleware developers have worked hand to hand with operation teams in order to test and validate the subsequent middleware versions introduced along the project. Only after a fully successful validation the new middleware has been installed at the level of the production

infrastructure. The price to pay has been investing an additional effort in order to maintain two testbeds, one for production and one for development, where the middleware has been previously tested and validated.

It is necessary to point out that the setup of a grid infrastructure supporting all these features requires an extra coordination effort because the integration of the components has to be specially tight. This is clearly the case when it comes to match application requirements that need to use resources in parallel, or request graphical and interactive access.

2 APPLICATIONS AND USE CASES

In the early phase of the project, we identified a relation of applications that benefit from the deployment of an infrastructure like *i2g*. This set of representative applications was analyzed and a detailed use case was worked out together with the middleware developers of the project. The complete set of requirements covered a reasonable spectrum of necessities concerning Grid implementation, interactive steering and visualization capabilities.

Evidently it is impossible to make a difference at the scientific level without the involvement of the researchers that are developing and using the application. Therefore we have analyzed, together with the code owners or the users, the hardware and software requirements, and how these requirements can be satisfied in a grid environment with the potential of gathering seamlessly a large amount of distributed resources.

The reference applications belong to four research areas: high energy physics, nuclear fusion, astrophysics, medical applications and environment. The details and full accounting of work done on some of them is reported in this volume as well. Let us nevertheless describe two examples to clarify further what is our working environment.

The first example is an application that visualizes the behavior of plasmas in nuclear fusion devices. There are several approaches to investigate the behavior of plasma particles in fusion reactors. In this application the plasma is analyzed as a many-body system, composed of electrons and ions, orbiting inside the fusion device cavities. This approach is becoming feasible only in the past two years because researchers can make profit of large scale clusters on their own centers, or distributed computing infrastructures like the Grid. Furthermore, the visualization of the plasma particles inside fusion devices is a very interesting tool because the nature of the plasma makes it very difficult to obtain direct experimental measurements in Tokamaks or Stellarators.

The visualization software has been developed by one *i2g* member institutes, BIFI, in collaboration with CIEMAT, as it reproduces the inside of the Stellarator TJ-II located in Madrid. In particular one can simulate the plasma behavior for an arbitrary configuration of goals. This code allows the visualization of the plasma particles orbiting from any angle of the reactor allowing researchers to develop a direct intuition about the particles behavior, especially because of the possibility of changing the parameters of the simulation in real time and see immediately what has been the effect on the plasma. This point is particularly important since Stellarators are extremely expensive to operate.

The purpose of porting this application to the Grid is exploiting the combined power of a large number of CPUs interactively on the Grid to visualize the behavior of plasmas inside fusion reactors, to be used as a designing tool for future fusion devices. In Figure 2 we show the application running inside the Migrating Desktop. This application uses most

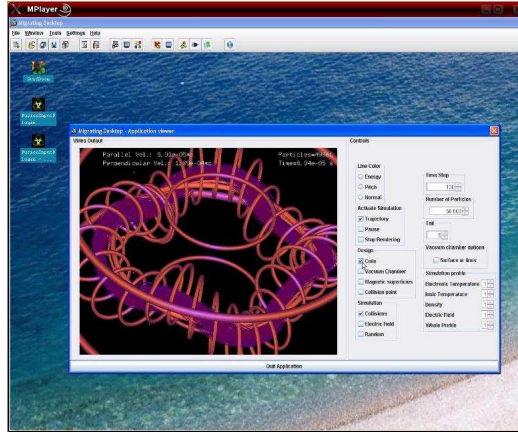


Fig. 2. Plasma visualization application running on the Migrating Desktop

of the capabilities of our middleware at the same time: MPI, interactivity and advanced visualization.

Another category of applications is medical imaging analysis. These are typically applications that handle large amounts of data that needs to be partitioned, analyzed and recombined again in order to see the effect of a particular filtering imaging technique, or in general software of image treatment. These applications come with extra constraints due to the nature of the users, which are usually researchers including doctors at hospitals. These are users without much experience in computing science, requiring web-based tools to handle job submission and retrieval, and having to fulfill legal conditions regarding patient privacy in data handling.

One of such applications is the Ultrasound Computertomography, USCT, a new method for medical imaging. It is based on the analysis of ultrasound data measured by a ultrasound scanner which surrounds the object of interest. The long term goal of the project is to develop a medical device for use in a hospital or at doctor's surgery, applicable for screening and location of breast cancer already at an early stage.

The current USCT prototype consists of 384 ultrasound emitters and 1536 receivers. In this setup every point (emitter) of a cylindrical scanner emits a pulse with a circular wave front, while all others receive the scattered signal. When the process is repeated for all emitters, it is possible to reconstruct the object that has produced the scattering patterns recorded by all transducers along the process. The researchers from the Forschungszentrum Karlsruhe are investigating the application of this technique for early detection of breast cancer. This means detection of tumors as small as possible, and with the most precise spatial location possible.

Conventional reconstruction techniques are slow and in practice impossible to use on a large scale in hospitals. The computational cost of image reconstruction is very high. In a typical situation we face an amount of data of about 20 GB. The reconstruction of the full data set at maximum resolution ($40\,962 \times 3\,410$) would take about 180 years in a single workstation. This would imply sub millimeter resolution for the whole volume. However,

such resolution is rarely needed, and in hospital cases one restricts the volume of study and also the areas of high resolution to be more limited. In a normal use case we would be talking about reconstruction times of about 30 days in a single workstation. Evidently this is too much time, and the goal is to reduce this time to something acceptable profiting from Grid technologies.

In general terms the CPU time needed for simulations in our project goes from 30 minutes in the medical application of Ultra Sound Computing Tomography (USCT) [9], up to 5–6 hours in an environmental application like the Pollution application [10]. The whole idea of interactivity and simulation steering relies on the fact that we are not dealing with runs of many hours, but rather with short intensive runs involving a large number of processors, and producing an output stream which the user can access in nearly real time. Simulation steering and runtime monitoring capability are some of the most important added values of the (*i2g*) middleware.

All (*i2g*) pilot applications require graphics display and visualization to some extent. The most demanding applications in these respect (Plasma Visualization, Evolution of Pollution Clouds, . . .) demand continuous access to the stream of data which reproduces the simulation on the screen. This data stream can be produced by a regular computing node or by a specialized one, a stream server, a specialized machine inside the eInfrastructure, which reads directly the output of the simulation and encodes it in video format. Such device can be of general interest for different research areas.

3 MIDDLEWARE, INFRASTRUCTURE AND INTEROPERABILITY

Interactive European Grid offers intra/inter cluster support for parallel applications, visualization and handling of video streams to support graphical applications, and all this with a mechanism to support on the fly response to user interaction.

A key development in application support has been the design and setup of *mpi-start*, a software layer that hides all the heterogeneity inherent to grid hardware and software setups (filesystem, MPI implementation,...) and allows for flexible and transparent use of MPI parallel applications. The consortium has chosen to adopt OpenMPI [11] as parallel library implementation because its modular approach makes it well suited for heterogeneous environments like the Grid [12, 13]. However, *mpi-start* can also handle MPICH, as is the case in EGEE, which has quickly adopted this solution for application clusters with demand of parallel support.

Support to interactivity and visualization is achieved using the *i2glogin* middleware together with the Gvid codec [14, 15]. The *i2glogin* middleware allows fully interactive connections to the Grid, with a functionality comparable to that of *ssh*, but without the overhead of a server side running on top of every Grid node. The Grid vid, or Gvid, codec has also been developed by the group at the University of Linz. It encodes the visualization in a video stream, saving bandwidth, and sends the output to any Grid resource, for example to User Interfaces. It supports mouse, keyboard and GUI events, which are transported over the network by *i2glogin*.

The *i2g* middleware relies on glite with additional enhancements. The Crossbroker [16] is a LCG enhanced resource broker. In particular it schedules and starts intra/inter cluster parallel jobs and handles video streams. The interaction of the user with the application has been made possible by the integration of *i2glogin* in the Resource Broker

of the *i2g* infrastructure, the Crossbroker [16]. There is also an improved User Interface with the possibility of submitting parallel jobs and handling visualization applications. Additionally, the computing element and the worker nodes contain *i2g* middleware to support all these enhancements.




Fig. 3. Middleware intersection with glite

The common part of glite middleware (see Figure 3) opens a window of interoperability with the rest of glite-based infrastructures across Europe. Indeed, interoperability has become an important point for Grid resource centers as described in the following section.

The *i2g* infrastructure [17] includes two independent testbeds. The production testbed is composed of clusters at nine different sites (see Figure 4) and intends to provide to users the quality of service they demand. Each cluster deploys different CPU architectures (Xeons, Opterons, Pentium Ds) and uses their preferred local resource management systems for job management control. This distributed production infrastructure represents more than 700 CPU cores and 16 TB of storage.

The development testbed aims at providing a realistic yet flexible environment making use of virtual machines. It includes four clusters and is used to support the testing and integration of the middleware and to allow early access to new middleware functionalities expected to become available in production.

As described in the previous section, the infrastructure supports several research communities. Currently there are 13 virtual organizations (VOs) with more than 200 users registered.

Regarding usage, more than 150.000 jobs have been executed, most of them in the last months, with a submission success rate above 90%. This is an acceptable rate given that many of them are yet development jobs requesting advanced features like MPI and interactivity usage, and we expect this number to improve once the final middleware release is consolidated.

Finally it must be said that the infrastructure has been set up with redundancy for almost all central critical services, with duplicated machines at LIP and IFCA sites, increasing their availability in a significant way.



Fig. 4. i2g infrastructure map

4 TOWARDS A SUSTAINABLE MODEL

Currently there are more than 25 Grid projects in Europe each serving different fields of science. In order to fulfill their objectives the projects deploy different middleware services. It is also frequent that the same users are in different projects depending on the applications requirements.

Interoperability allows a common layer which will help users access different grid infrastructures more transparently.

Moreover, when considering the exploitation and sustainability of this e-Infrastructure oriented to support interactive applications, interoperability is the key when considering from a site administration point of view how to integrate these resources: it avoids the duplication of the effort on installation and maintenance.

Interoperability also allows to setup a framework where the different actors that include resource owners (infrastructure sites), and users (research communities) can establish agreements and negotiate new conditions without being constrained by the requirement of a very specific service or middleware.

Finally, in the case of *i2g*, the Crossbroker glide-in mechanism allows that the very same resources can be shared both for interactive and batch jobs, optimizing the usage of an infrastructure that is available 24h, but where the main interactive load concentrates along working hours.

An exploitation phase of the whole e-Infrastructure will be explored in the coming months, and in particular it is expected that agreements with new EU projects to support research communities, like EUFORIA or DORII, will benefit from our project.

5 CONCLUSIONS

The Interactive European Grid (*i2g*) project has proved that researchers can benefit from access to powerful resources in interactive mode for the final stage of their analysis, including the use of parallelism, visualization, and within a user friendly framework. In particular advanced applications in research areas like physics and medicine have been deployed in the (*i2g*) infrastructure, and exploiting the possibilities provided by our middleware, not available in other frameworks. The middleware developed and the infrastructure setup adopted assure interoperability with largest existing e-Infrastructures in Europe, assuring good perspectives for its sustainability, and we look forward to a fruitful exploitation phase.

REFERENCES


- [1] The Grid. Blueprint for a New Computing Infrastructure. I. Foster and C. Kesselmann (Eds.), Morgan Kaufmann 1999.
- [2] See official web page of the EGEE project - <http://www.eu-egee.org>.
- [3] See official web page of the OSG project - <http://www.opensciencegrid.org>.
- [4] Background and achievements of the EU DataGrid project are available at <http://eu-datagrid.web.cern.ch>.
- [5] The official page of glite is <http://cern.ch/glite>.
- [6] Information about the Globus Alliance can be found at <http://www.globus.org>.
- [7] GOMES, J. et al.: Experience with the International Testbed in the CrossGrid Project. CrossGrid Consortium, Proceedings of the European Grid Conference 2005 (EGC 2005), Science Park, Amsterdam, The Netherlands, February 14–16, 2005.
- [8] KUPCZYK, M.—LICHWALA, R.—MEYER, N.—PALAK, B.—PŁÓCIENNIK, M.—WOLNIEWICZ, P.: Applications on Demand as the Exploitation of the Migrating Desktop. Future Generation Computer Systems, Vol. 21, 2005, No. 1, pp. 37–44.
- [9] YARKHAN, A.—SEYMOUR, K.—SAGI, K.—SHI, Z.—DONGARRA, J.: Recent Developments in GridSolve. International Journal of High Performance Computing Applications (Special issue: Scheduling for Large-Scale Heterogeneous Platforms, Y. Robert (Ed.)), Vol. 20, 2006, No. 1.
- [10] ŠIMO, B. et al. (see this issue).
- [11] GABRIEL, E. et al.: Open MPI: Goals, Concept and Design of a Next Generation MPI Implementation. Proceedings of 11th European PVM/MPI Users' Group Meeting, Elsevier 2004, pp. 97–104.
- [12] DICHEV, K.—KELLER, R.—FERNÁNDEZ, E. (see this issue).
- [13] FERNÁNDEZ, E.—HEYMANN, E.—SENAR, M. A.: Suporting Effiicent Execution of MPI Applications Accross Multiple Sites. 12th international Euro-Par Conference, Dresden (Germany), LNCS, Vol. 4128, 2006, pp. 383–392.
- [14] FERNÁNDEZ, E.—HEYMANN, E.—SENAR, M. A.: Proceedings of IEEE Int. Conf. on Cluster Computing (Cluster 2006), Barcelona (Spain), CD-ROM edition, IEEE CS Press, September 2006.

- [15] ROSMANITH, H.—VOLKERT, J. (see this issue).
 [16] FERNÁNDEZ, E.—CENCERRADO, A.—HEYMANN, E.—SENAR, M. A. (see this issue).
 [17] GOMES, J. et al. (see this issue).



Jesús MARCO DE LUCAS is CSIC Research professor. He is responsible for the distributed computing research line. Ph. D. in particle physics in 1989 working for the DELPHI collaboration at CERN, where he became later coordinator of the Higgs Research Line. He has published over 300 papers in the HEP field. Since 2001 he has promoted and participated in several Grid projects, in particular he has acted as coordinator of the international testbed of the EU CrossGrid project (2002-2005), and leader of the integration team. He has published several contributions in this field, and made presentations at relevant forums (including

GGF, ISTconcertation meetings, AcrossGrid conferences, etc.). He promoted the IrisGrid initiative in Spain, was member of the team that prepared the White Book for e-Science in Spain, and has participated in the e-IRG forum. He is the coordinator of the Interactive European Grid project.



Robert PAJAK received his M. Sc. degree in computer science from the University of Mining and Metallurgy (AGH), Krakow, Poland, in 2001. He was participating in several Polish and international projects in the high performance computing area, including CrossGrid, EGEE and K-Wf Grid, working mainly in dissemination and education area. Now he is involved in Baltic-Grid and GREDIA projects. He is also a leader of the Dissemination and Outreach work package of the Int.eu.grid project.



Isabel Campos PLASCENCIA is the Staff Scientist of CSIC at the Instituto de Física de Cantabria (IFCA) in Santander (Spain). Ph. D. in Lattice Quantum Chromodynamics (QCD) in 1998 at the University of Zaragoza (Spain). Since then she held research associate positions at the Deutsches Elektronen Synchrotron in Hamburg and Brookhaven National Laboratory in New York. After that she worked in the group of applications and development on the Linux Cluster at the Leibniz Computer Center in Munich from April 2002 till January 2004, when she joined the institute of Biocomputing and Physics of Complex Systems

in the University of Zaragoza (Spain). In August 2006 she joined IFCA as applications coordinator in the group of Distributed Computing lead by Jess Marco. She has 27 publications in international peer reviewed journals and over 30 proceedings in presentations to international conferences.



Jorge Gomes is a senior computing research scientist at LIP where he coordinates the Computer Centre and the grid computing team. He is currently participating in the EGEE, i2g, LCG and IBERGRID projects. He is the i2g infrastructure operations manager.



Miquel Angel Senar is Associate Professor of Computer Architecture at the University Autònoma of Barcelona (Spain). He attained his B.Sc. degree in 1988, and his Ph.D. in 1996 at the same university. Throughout his career he has conducted research and published extensively in the areas of computer-based applications, parallel and distributed processing and high performance architectures. His current research interests include resource management on distributed systems, tool interoperability in distributed systems and scheduling of applications in parallel and distributed systems. He has participated in several research

projects related with the development of middleware tools for parallel and distributed systems.