EUFORIA: Grid and High Performance Computing at the Service of Fusion Modelling

Francisco Castejón[1], Miguel Cárdenas-Montes[2], Antonio Gómez-Iglesias[2], Enrique Morales-Ramos[2], Bernard Guillerminet[3], David P. Coster[4], Eric Sonnendrücker[5], Isabel Campos Plasencia[6], Jan Åström[7], Jan Westerholm[8], José María Cela[9], Leon Kos[10], Lorna Smith[11], Marcin Plociennik[12], Marcus Hardt[13], Mats Aspnäs[8], Pär Strand[14], Rainer Stotzka[13]

[1] CIEMAT, Madrid, Spain
`francisco.castejon@ciemat.es`
[2] Extermadura Advanced Research Center (CETA-CIEMAT), Trujillo, Spain
{`miguel.cardenas, antonio.gomez, enrique.morales`}`@ciemat.es`
[3] Commissariat à l'énergie atomique (CEA), France
`Bernard.Guillerminet@cea.fr`
[4] Max Planck Institute for Plasma Physics (IPP), Germany
`dpc@ipp-garching.mpg.de`
[5] Université Louis Pasteur, Strasbourg, France
`sonnen@math.u-strasbg.fr`
[6] Instituto de Física de Cantabria (IFCA), CSIC, Santander, Spain
`iscampos@ifca.unican.es`
[7] Finnish IT Center for Science (CSC), Finland
`jan.astrom@csc.fi`
[8] Åbo Akademi University (ABO), Finland
{`jan.westerholm, mats.aspnas`}`@abo.fi`
[9] Barcelona Supercomputing Center (BSC), Barcelona, Spain
`cela@ac.upc.es`
[10] University of Ljubljana (LECAD), SI-1000, Ljubljana, Slovenia
`leon.kos@lecad.uni-lj.si`
[11] University of Edinburgh (EPCC), Edinburgh, United Kingdom
`l.smith@epcc.ed.ac.uk`
[12] Poznan Supercomputing and Networking Center (PSNC), Poznan, Poland
`marcinp@man.poznan.pl`
[13] Forschungszentrum Karlsruhe (FZK), Karlsruhe, Germany
`Marcus.Hardt@iwr.fzk.de, rainer.stotzka@ipe.fzk.de`
[14] Chalmers University of Technology (CHALMERS), Göteborg, Sweden
`elfps@chalmers.se`

Abstract. ITER is the next generation of fusion devices and is intended to demonstrate the scientific and technical feasibility of fusion as a sustainable energy source for the future. To exploit the full potential of the device and to guarantee optimal operation for the device a high degree of physics modelling and simulation is needed already in the current construction phase of the ITER project. First principles modelling tools that are needed for an adequate description of the underlying physics cover both a wide range of timescales and spatial orderings and are in general very demanding from a computational point of view. An improved access to computing infrastructures will be instrumental in advancing a pan-European modelling activity for ITER to a very competitive status in relation to the ITER partners. In this paper the main goals and features of the EUFORIA, a grid and high performance computing project in fusion modelling, are presented.

Keywords: Fusion, HPC, Grid Computing, ITER.

1 Introduction

The EUFORIA project will provide a comprehensive framework and infrastructure for core and edge transport and turbulence simulation, linking Grid and High Performance Computing (HPC), to the fusion modelling community. The project will enhance the modelling capabilities for ITER and DEMO[1] sized plasmas[2] through the adaptation, optimisation and integration of a set of applications for edge and core transport modelling targeting different computing paradigms as needed (serial and parallel grid computing[3] and HPC). As in contemporary and related projects, like PRACE (the Partnership for Advanced Computing in Europe)[4], the deployment of both a Grid and a High Performance Computing services are essential to the project. A novel aspect is the dynamic coupling and integration of codes and applications running on a set of heterogeneous platforms into a single coupled framework through a workflow engine, a mechanism needed to provide the necessary level integration in the physics applications. This strongly enhances the integrated modelling capabilities of fusion plasmas[5] [6] [7] and will at the same time provide new computing infrastructure and tools to the fusion community in general. Figure 1 shows the way in which EUFORIA project would evolve:

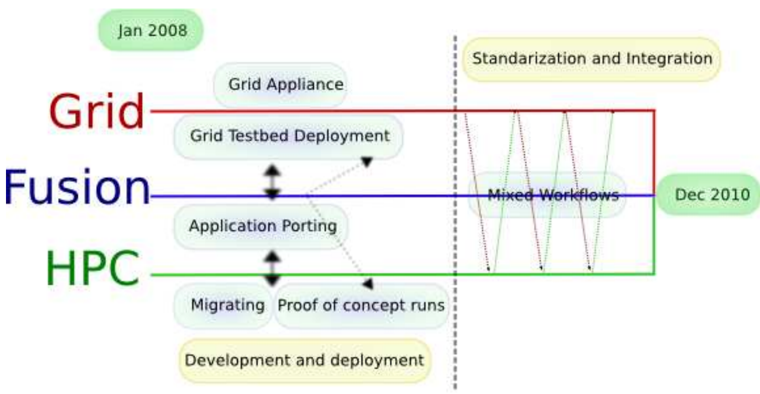

Fig. 1. The predicted roadmap for EUFORIA project.

This paper consists of the following sections: section 2 presents an overview of the core technologies involved in the project. In section 3, the work in the adaptation of codes for the project's infrastructure is described. Then, in sections 4 and 5 the work lines in workflow orchestration and data visualization are described. Finally, section 6 presents the main conclusions.

2 Core Technologies and Infrastructures

Within the fusion community, a large number of codes are in use to simulate various aspects of the plasma behaviour. With very few exceptions, these codes have been written by physicists with an emphasis on the physics, with a much smaller emphasis on using the latest technologies available from the computer science community[8]. The EUFORIA project aims to improve the situation in a number of ways like increasing the performance of key existing codes and enabling a subset of the existing codes for use on computational grids[9].

In order to focus the effort, an important physics sub-domain was identified, that of **transport and turbulence in the core and edge**. Within this area, a subset of the available codes has been identified.

By taking a representative set of codes, it is hoped that the "lessons learned" will be applicable to the much wider set of codes that are being used in the chosen physics area, and, in fact, in other plasma physics areas as well.

In order to simulate the behaviour of the plasma in a tokamak, codes have been written to address each of the sub-problems identified by one of the boxes. Part of the EUFORIA project will concentrate on improving the performance of each of the individual codes, and preparing them for the computational resources of the future. Another part of the project concentrates on facilitating the inter-connections between the boxes (the "arrows" in the diagram), and the remaining part concentrates on methods of understanding the results by implementing and improving new visualization techniques. Figure 2 summarizes part of the physics problem that this proposal addresses:

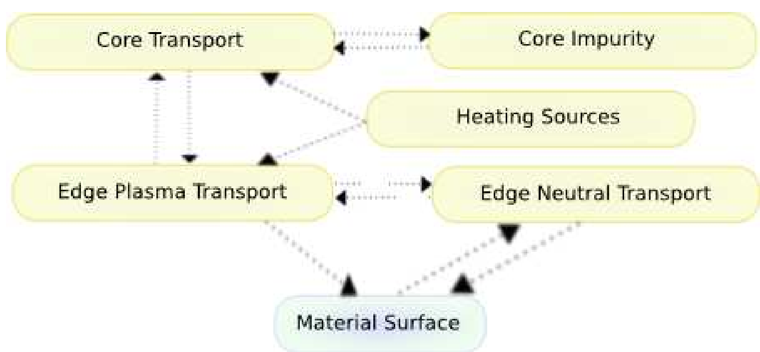

Fig. 2. A simplified figure showing some of the inter-relationships involved in simulating fusion plasmas.

2.1 Infrastructure Needs

As an heterogeneous range of computational resources are required, this approach relies heavily on the existence and guaranteed availability of the needed computational infrastructures[10] and the organisation and transfer of data between applications. To support this structured integration of codes, a workflow orchestration tool is envisaged for scheduling applications, transferring data, integrating results and visualisation tools transparently for the end user. A common ontology for the edge and core simulations will be developed based on ongoing European developments towards a mark up language for fusion simulations, data access and storage. Access to the EUFORIA core-edge simulation tool will be through three-tiered approach where an access portal is used to launch the workflow orchestration tool and access to data where simulation are run on distributed set of resources ranging from grid services to HPC computing platforms.

An important issue for the integration and utilisation of a whole range of different codes is to provide tools for natural and efficient processing and visualisation of the generated data in order to be able to interpret the results and compare with experimental data. In addition to codes generating 1D or 2D data for which there are now standard visualisation tools available, many of the codes involved generate massive and complex 3D data or even 4D or 5D data for the distribution function computed in gyrokinetic codes. Hence visualisation support will rely on at least two aspects. First an unique compressed data format that can be handled efficiently needs to be defined, then a single visualisation environment should to be included both in the workflow to enable monitoring of the simulation and for post-processing of the compressed data that have been saved during the simulation. This environment should be able to produce in an easy and natural manner line plots, contour plots and other simple plots and also enable the user to generate more sophisticated specific plots when needed.

2.2 Implementation

The EUFORIA project consists of two different phases that are partly being developed in parallel from the start of the project to become fully integrated in the later stages.

The first is a development and deployment phase consisting of the adaptation and optimisation of a selection of codes covering both Edge and Core Physics for Grid and HPC environments as appropriate. Inherent to this activity is the deployment of the computational infrastructure where access to grid computing and high performance computing is being developed for the project members.

The second phase is a standardisation and integration activity and has one technology driven part which develops the technologies and tools to provide user transparent methods for resource allocation and scheduling and dynamic coupling of physics codes and one user and physics driven component where a technology for building complex workflows with the optimised codes as components and standardised data structures and transfer methods are being utilised to expand the physics cases for ITER development.

3 Adaptation of Codes for HPC and HTC

The project puts a strong focus on the infrastructure operation, aggregating all related tasks under a single management structure of service activities that will strongly contribute to improve the coordination of the operational efforts. Especially appealing from a research and development point of view are the activities dedicated to make joint use of Grid and HPC Infrastructures. There are a number of codes which can run on both facilities depending on the amount of resources needed. Situations, for example, in which a parallel code needs to run on less than 64 processors to analyse the basic physical behaviour, and afterwards, needs to run on more than 256 processors in order to refine measurements or check scaling behaviour is very likely to arise.

Bearing in mind another analogous research initiatives[4] and as part of the preparatory work for the EUFORIA project contact was made with the European Fusion Laboratories and University based research groups carrying activities in the core and edge transport areas. A request asking them to express their interest in committing codes and simulation tools for EUFORIA was presented and the responses analysed and catalogued. The strong interest from the user and code developer community has been one of the driving forces behind the creation of EUFORIA as there is a very clear need to bridge limited national resources and provide more coherent approach to the deployment and utilization of these tools[11].

The following codes have been identified as of potential interest for their adaptation:

BIT1 Kinetic 1D3V (1D in usual and 3D in velocity space) code for simulation of the plasma edge.
CENTORI Fully toroidal two-fluid, electromagnetic turbulence simulation code.
COREDIV Transport of energy code, main ions and impurity ions in the core and the scrape of layer regions.
EIRENE Kinetic neutral particle and line radiation transport code.
ELMFIRE Gyro-kinetic full-f particle code, with mostly global emphasis.
ERO Gyro-kinetic code for impurity transport in plasma.
ESEL Turbulence and profile evolution code at the outboard midplane in the SOL, using a fluid (ESEL) and gyrofluid (GESEL) approach.
GEM Gyrofluid code (GEM is local, GEMX is nonlocal, same infrastructure, similar scheme and programming, somewhat differently formulated equations).
GENE Nonlinear gyrokinetic code to investigate plasma turbulence.
ISDEP Kinetic theory of transport code based on Langevin Equations[12].
SOLPS Two codes (B2-Eirene) tightly coupled together based code.
TECXY 2D multifluid plasma and impurity transport in the tokamak edge simulating code.
TYR Drift Alfven plasma fluid turbulence and transport in flux-tube geometry code.

The currently identified Grid-serial codes are BIT1, COREDIV, EIRENE, ERO, ESEL, ISDEP and SOLPS. Except for EIRENE, the rest of listed codes are being also considered for HPC parallelization. It is also possible and likely that additional codes might be identified. Of the above list of Grid-serial codes, all are also interested in Grid-parallel versions as well. In addition, the following codes are only of Grid interest in Grid-parallel versions: CENTORI, ELMFIRE, GEM, GENE, TECXY and TYR. All of them are also going to be studied and becoming candidates to be HPC ported.

3.1 Adaptation of Codes for Grid Infrastructure

The consortium plans to deploy a grid infrastructure shared between all the members dedicated to Fusion Research. To achieve its mission a significant amount of high value computing resources will be made available to the fusion researchers as a grid infrastructure.

The main objective of this activity is the operation and integration of these resources with a production-level performance as a Grid-empowered e-Infrastructure.

Joining the resources of a particular site to a global grid infrastructure requires the dedication of a number of service machines. Evidently this calls for the use of virtualisation techniques[10] to optimise resources. Therefore one of the milestones of the project is to design, deploy and maintain a single server running all the services needed by a local site to connect the resources to the Grid.

Grid applications will be analysed and separated into categories according to computing necessities (MPI applications and serial applications, RAM consumption, input/output size, etc...).

Operation and access to the Grid-empowered infrastructure The operational scheme for the e-Infrastructure is a key point. A simplified operational division, based on the logical roles and administrative domains of the actors involved, can be considered initially:

- **Operation of Computing Resources:** The computers are installed and administered locally at each centre; their integration in the Grid framework requires the installation of services oriented to provide access.
- **Operation of Virtual Organisations:** The Virtual Organisation should provide a framework to support and manage the user's activity. To access the resources the user must first get authenticated. Next the user should contact the corresponding VO administrator for authorisation.

Mechanisms and service-oriented tasks For the purpose of organisation a hierarchy of responsibility and coordination must exist. This hierarchy is composed of two layers:

1. The top layer contains a single infrastructure management body named **Grid Operations Management** (GOM) that is responsible for the management and coordination of the infrastructure operations.
2. The **Resource Providers** (RP) populate the bottom layer: these are the sites that offer resources to the infrastructure.

Figure 3 presents the mentioned hierarchy:

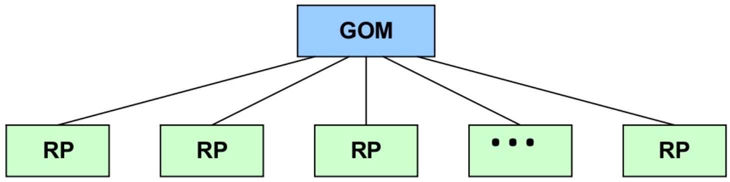

Fig. 3. Schematic view of the Grid infrastructure.

The GOM will coordinate all aspects of the infrastructure operation. It will interact with other project activities to coordinate the Grid operation actions, to provide or to request services and to provide or to obtain feedback to improve the infrastructure and operations quality. The GOM coordinates the operation services needed to address the specific aspects of the operations activity.

1. **Accounting:** Deploys accounting systems to obtain information about the usage of the infrastructure resources such as CPU and storage. The information provided by the accounting will be very important to evaluate the infrastructure capacity and success.
2. **Helpdesk:** Provides infrastructure support to VOs and resource providers to address issues related with authentication, authorization, VO setup, Grid services and Grid resources.
3. **VO management:** Is responsible for the virtual organizations management aspects such as the provisioning of the required services for the VOs and the configuration of the infrastructure to support the VOs.
4. **Security management:** Keeps track of software vulnerabilities, and make sure that the recommended corrections are applied.
5. **Central services:** Deploys, operates and maintains the core services necessary to integrate the infrastructure resources into a coherent homogeneous Grid service.
6. **Monitoring:** Deploys, operates and maintains monitoring services aimed to obtain accurate information about the infrastructure status.

7. **Deployment coordination:** Coordinates the deployment and maintenance the middleware releases. This is one of the most important parts of the operations since the stability of the infrastructure will depend much on the efficiency of the infrastructure deployment and maintenance.

Figure 4 illustrates the structure of the GOM services that have been described:

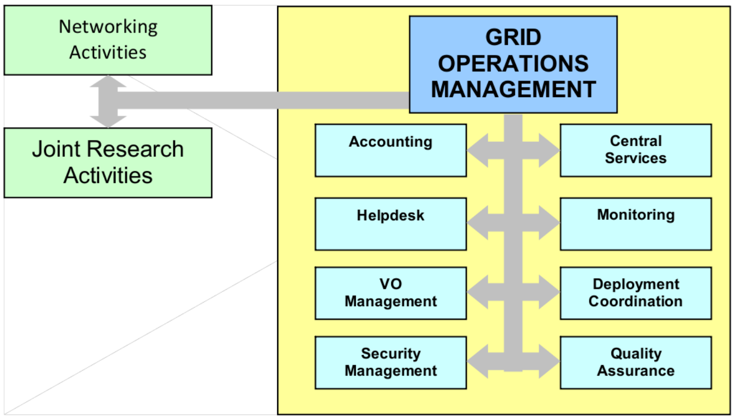

Fig. 4. Schematic representation of the structure of services.

3.2 Adaptation of Codes for HPC Infrastructure

In Computational Plasma Physics there are several computational schemes that require fast communication between computer processors. Examples include fluid theory, gyro-kinetics, turbulence, heating and current drive. Massively Parallel supercomputers (MP-SC) with fast communication between processors are the only possible tools for efficient investigation of this type of phenomena. In Europe, the type of MP-SCs that is needed for this task is typically found at a small number of large national supercomputer centres.

The code developers are focused on optimising codes from within the plasma fusion community that need MP-SCs to run efficiently. In order for the optimisation work to function properly the code developers must have access to suitable computers for testing, debugging, and running.

Programming Environment The programming environment includes e.g. licence to use compilers, debuggers and, if available, software profiling tools. The aim will be to ensure that EUFORIA users have the appropriate tools to utilise these HPC resources. Access to run time on MP-SCs is provided, free of charge, to such an extent that can be motivated for testing of the optimized codes.

Proof of Concept This activity will provide a relatively small amount of HPC resource to users to demonstrate the effectiveness of the codes and their potential to carry out world leading science. This is essential to demonstrate the success of the adaptation of codes for the HPC infrastructure and will require a relatively small amount of CPU time. This service cost will provide CPU time, data storage capacity, etc. included in MP-SC infrastructures and staff effort from the main support staff of the HPC systems.

World Leading Calculations The next step will utilise the results of the proof of concept work to apply for larger amounts of time on relevant HPC resources at zero cost, for the community to carry out world leading simulations. This time will be obtained through applications to national resource committees and through the DEISA Extreme Computing Initiative.

HPC Resources Between the involved supercomputing centres, EUFORIA users will have access to an extensive range of leading edge HPC resources. These resources include massively parallel distributed systems and clustered Symmetric Multi Process systems, providing enough diversity to match the requirements of the range of different application codes.

4 Workflow Orchestration

The main objective of this activity is to schedule jobs on the GRID and HPC infrastructures together with jobs running on others computing facilities. Jobs must be launched and controlled in a transparent manner for the users. Data communication is part of this activity since data transfers and visualisations are required at run time and to access the experimental and simulated data. The tasks involved in this activity are to develop:

- a Java library to address the gLite middleware[13] within the project
- the necessary tools to notify the external scheduler of the job status
- the data communication layer for on-line visualisation and intra task communication
- the data communication layer to communicate with the Universal Access Layer (UAL)[14]
- a tool which schedules the jobs on the different infrastructures using the workflow described in MoML[15]

5 Visualisation

In order to perform an end-to-end numerical simulation of a Tokamak discharge, several complex codes will need to be coupled. In order to provide easy and efficient access to these tools and the ability to efficiently make many runs with different parameters they will be integrated as actors in the scientific workflow package Kepler[16] which is Java based. There will also be the need for unified visualisation support[17] so that scientists can use it without having to learn to use specific tools for each code that has been integrated. Figure 5 shows a running example of KEPLER tool:

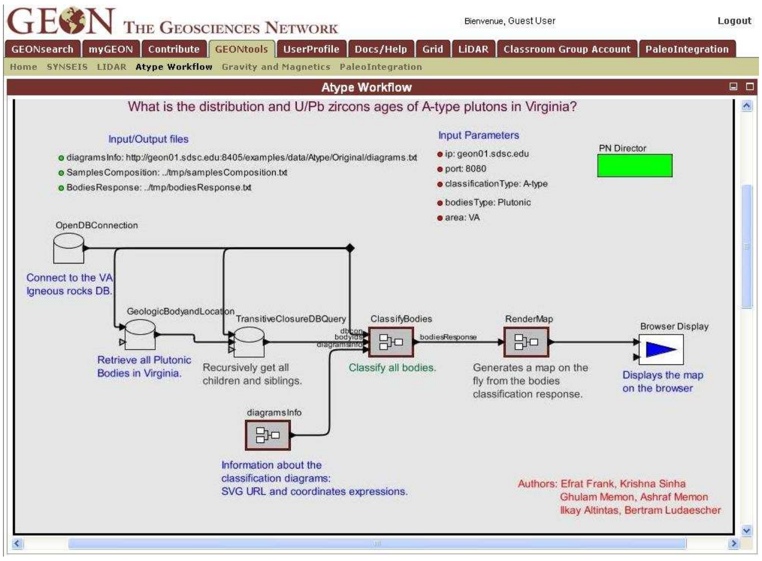

Fig. 5. KEPLER: workflow design tool in GEON portal.

This visualisation support will consist of two aspects. On the one hand some specific visualisation tools must directly be inserted in the Kepler workflow in order to be able to monitor ongoing simulations. On the other hand, data generated by simulations must be thoroughly analysed in a post-processing step. In this case the coupling with the visualisation tools only consists in enabling them to read the data formats generated by the codes.

6 Conclusions

In a similar way as other new-born projects[4], the EUFORIA project attempts to offer more and more useful computing and storage power for scientific aims and challenges. It will improve the modelling capacities for the new fusion device called ITER. These improvements are mandatory to exploit the full potential of the device arising the physical modelling skill. There are a number of Plasma Physics problems that can be solved by Grid computing. Nevertheless, applications that run on the Grid can run efficiently in supercomputers, although the opposite is not necessarily true. The future Plasma Physics computations will involve mixed problems with the combination of Grids and HPCs and more complex workflows are expected: Rather than thinking of programs or codes, it will be necessary to think of the tools as processes that are launching other subprocesses that can run on different architectures. The main topics presented in this paper include the selected codes to be ported, the develoment and deployment needed to ease the simulation and analysis and, finally, the work in workflow orchestration and data visualization which will improve the results and make them easier to use.

References

1. International Thermonuclear Experimental Reactor. Website. http://www.iter.org/ February 2008.
2. L. Jones, J. Duhovnik, M. Ginola, J. Huttunen, K. Ioki, L. Junek, T. Loewer, U. Luconi, M. Pick, G. -P. Sanguinetti, M. Slovacek, Y. Utin. Results from ITER vacuum vessel sector manufacturing development in Europe. FUSION ENGINEERING AND DESIGN. Pages 1942-1947. October 2007.
3. Antonio Gómez-Iglesias, Miguel A. Vega-Rodríguez, Francisco Castejón Magaña, Manuel Rubio del Solar, Miguel Cárdenas Montes. Optimizing the configuration of magnetic confinement devices with evolutionary algorithms and grid computing. Third EELA Conference. Catania, Italy, pp. 91-99. December 2007.
4. Partnership for Advanced Computing in Europe (PRACE). Website. http://www.prace-project.eu/ February 2008.
5. N. Jelic, K.-U. Riemann, T. Gyergyek, S. Kuhn, M. Stanojevic, J. Duhovnik. Fluid and kinetic parameters near the plasma-sheath boundary for finite Debye lengths. PHYSICS OF PLASMAS. October 2007.
6. M. Stanojevic, J. Duhovnik, N. Jelic, A. Kendl, S. Kuhn. Fluid model of the magnetic presheath in a turbulent plasma. PLASMA PHYSICS AND CONTROLLED FUSION. Pages 685-712. May 2005.
7. Mladen Stanojevic, Joze Duhovnik, Nikola Jelic, Siegbert Kuhn. Magnetic presheath in a weakly turbulent multicomponent plasma. PHYSICS OF PLASMAS. January 2007.
8. Ian Foster, Carl Kesselman, Steven Tuecke. The Anatomy of the Grid. Enabling Scalable Virtual Organizations. International Journal of Supercomputing Applications. 2001.
9. F. Castejón, I. Campos, A Cappa, L. A. Fernández, E. Huedo, I. M. Llorente, V. Martín, R. S. Montero, M. Mikhailov, A. Tarancón, M. A. Tereschenko, J. L. Vázquez-Poletti, J. L. Velasco, V. Voznesensky. Fusion in the Grid. Spanish Conference on e-Science Grid Computing. Madrid, Spain, pp. 141-146. March 2007.
10. Miguel Cárdenas Montes, Javier Pérez-Griffo Callejón, Manuel Rubio de Solar, Raúl Ramos Pollán. Management of a grid infrastructure in GLITE with Virtualization. Ibergrid'2007 First Iberian Grid Infrastructure Conference, Santiago de Compostela, Spain, pp. 313-321, May 2007.
11. F. Castejón, J.L. Velasco, A. Tarancón, J. L. Vázquez-Poletti, V. Voznesensky, M. Tereshchenko, A. Cappa. Fusion Results within EGEE. EGEE 3rd User Forum, Clermont-Ferrand (France). February 2008.

12. I. Campos, F. Castejón, G. Losilla, J. M. Reynolds, F. Serrano, A. Tarancón, R. Vallés, J. L. Velasco. IVISDEP: a Fusion plasma application ported to the Interactive European Grid e-Infrastructure. Spanish Conference on e-Science Grid Computing. Madrid, Spain, pp. 129-133. March 2007.
13. gLite Middleware. Website. http://glite.web.cern.ch/glite/ February 2008.
14. EFDA Task Force ITM. Website. http://www.efda-taskforce-itm.org/ February 2008.
15. Edward A. Lee, Steve Neuendorffer. MoML - A Modeling Markup Language in XML - Version 0.4. Technical Memorandum ERL/UCB M 00/12. March 2000.
16. Kepler Project. Website. http://kepler-project.org/ February 2008.
17. Herbert Rosmanith, Jens Volkert, Rubén Vallés, Fermín Serrano, Marcin Plociennik, Michal Owsiak. Interactive Fusion Simulation and Visualisation on the Grid. Parallel and Distributed Computing, 2007. ISPDC '07. July 2007.