

Dipartimento di Informatica  
Università del Piemonte Orientale “A. Avogadro”  
Via Bellini 25/G, 15100 Alessandria  
<http://www.di.unipmn.it>



**The ShareGrid Portal: an easy way to submit jobs on computational Grids**

*Authors: Cosimo Anglano, Massimo Canonico and Marco Guazzone  
([cosimo.anglano,massimo.canonico,marco.guazzone](mailto:cosimo.anglano,massimo.canonico,marco.guazzone@unipmn.it))@unipmn.it)*

TECHNICAL REPORT TR-INF-2008-10-08-UNIPMN  
(October 2008)

The University of Piemonte Orientale Department of Computer Science Research Technical Reports are available via WWW at URL  
<http://www.di.unipmn.it/>.  
Plain-text abstracts organized by year are available in the directory

## Recent Titles from the TR-INF-UNIPMN Technical Report Series

- 2008-07 *BuzzChecker: Exploiting the Web to Better Understand Society*, Furini, M., Montangero, S., July 2008.
- 2008-06 *Low-Memory Adaptive Prefix Coding*, Gagie, T., Nekrich, Y., July 2008.
- 2008-05 *Non deterministic Repairable Fault Trees for computing optimal repair strategy*, Beccuti, M., Codetta-Raiteri, D., Franceschinis, G., July 2008.
- 2008-04 *Reliability and QoS Analysis of the Italian GARR network*, Bobbio, A., Terruggia, R., June 2008.
- 2008-03 *Mean Field Methods in performance analysis*, Gribaudo, M., Telek, M., Bobbio, A., March 2008.
- 2008-02 *Move-to-Front, Distance Coding, and Inversion Frequencies Revisited*, Gagie, T., Manzini, G., March 2008.
- 2008-01 *Space-Conscious Data Indexing and Compression in a Streaming Model*, Ferragina, P., Gagie, T., Manzini, G., February 2008.
- 2007-05 *Scheduling Algorithms for Multiple Bag-of-Task Applications on Desktop Grids: a Knowledge-Free Approach*, Canonico, M., Anglano, C., December 2007.
- 2007-04 *Verifying the Conformance of Agents with Multiparty Protocols*, Giordano, L., Martelli, A., November 2007.
- 2007-03 *A fuzzy approach to similarity in Case-Based Reasoning suitable to SQL implementation*, Portinale, L., Montani, S., October 2007.
- 2007-02 *Space-conscious compression*, Gagie, T., Manzini, G., June 2007.
- 2007-01 *Markov Decision Petri Net and Markov Decision Well-formed Net Formalisms*, Beccuti, M., Franceschinis, G., Haddad, S., February 2007.
- 2006-03 *New challenges in network reliability analysis*, Bobbio, A., Ferraris, C., Terruggia, R., November 2006.
- 2006-03 *The Engineering of a Compression Boosting Library: Theory vs Practice in BWT compression*, Ferragina, P., Giancarlo, R., Manzini, G., June 2006.
- 2006-02 *A Case-Based Architecture for Temporal Abstraction Configuration and Processing*, Portinale, L., Montani, S., Bottrighi, A., Leonardi, G., Juarez, J., May 2006.
- 2006-01 *The Draw-Net Modeling System: a framework for the design and the solution of single-formalism and multi-formalism models*, Gribaudo, M., Codetta-Raiteri, D., Franceschinis, G., January 2006.
- 2005-06 *Compressing and Searching XML Data Via Two Zips*, Ferragina, P., Luccio, F., Manzini, G., Muthukrishnan, S., December 2005.
- 2005-05 *Policy Based Anonymous Channel*, Egidi, L., Porcelli, G., November 2005.
- 2005-04 *An Audio-Video Summarization Scheme Based on Audio and Video Analysis*, Furini, M., Ghini, V., October 2005.

# The ShareGrid Portal: an easy way to submit jobs on computational Grids \*

Cosimo Anglano

Massimo Canonico

Marco Guazzone

Department of Computer Science, University of Piemonte Orientale, Alessandria (Italy),

email: {cosimo.anglano, massimo.canonico, marco.guazzone}@unipmn.it

## Abstract

*Grid computing is a distributed computing paradigm which aims to aggregate several heterogeneous and distributed resources, belonging to different and independent organizations, in a dynamic, transparent and coordinated way. Since its introduction, Grid computing has been successfully applied to solve several scientific challenging applications. Despite of the consolidation of many of its aspects, there are some issues that are still open. One of them is the transparency: in many real Grid systems, users still need to be aware of Grid computing, either for adapting their applications to this paradigm or for wrapping them in a suitable software framework. In this paper we present the ShareGrid Portal, a Web portal and a portal framework, built on top of the ShareGrid project infrastructure. Its intent is both to ease the execution of user applications in a Grid system and to allow developers to flexibly add new portal functionalities. In this work, we compare it with other well-known Grid portals and we show its user interface and its architecture. Finally we discuss user experiences and future extensions.*

**Keywords:** Grid computing, Grid portal, Portal framework, Bag-of-Tasks, Parameter sweep applications.

## 1 The ShareGrid Portal

### 1.1 Introduction

With the advances of technology used for building scientific instrumentation, scientists are now able to explore aspects that only few years ago were not completely able to observe. As a consequence, scientific applications, used to carry on scientific experiments, might need to access and elaborate massive amounts of data at a faster rate, and,

therefore, they are characterized by an high demand of storage and computational resources. For instance, the *Large Hadron Collider* (LHC) [5], the world's largest and highest-energy particle accelerator built by *CERN* [2], is able to produce petabytes of data per year that need to be accessed and studied by thousands of scientists belonging to different organizations and spread around the world. The computing paradigm traditionally used for running compute-intensive applications is the *cluster computing* [15], which implies the use of a powerful cluster of homogeneous computers owned by a single organization. This paradigm does not fit very well with the above scenario, since the available resources might be insufficient to satisfy the incoming demand of computational power and storage space, and it generally does not allow collaboration among users of different organizations. Hence, a more complex computing paradigm is needed in order to allow multi-institutional collaboration and resource sharing in a way as easy as possible. The Grid computing [21] is a computing paradigm that tries to achieve this goal through a software and hardware infrastructure in order to allow and ease the dynamic sharing of heterogeneous resources among different, independent and geographically distributed organizations, in a way that should be totally transparent to the user.

The coordination and the sharing of heterogeneous resources between different communities distributed on a large geographic scale, has partially changed the way of solving compute-intensive and data-intensive applications. Being a kind of a distributed system, one of the initial promise of the Grid computing was to make the execution of applications in a Grid system completely transparent from the point of view of users: the user of such a system should be unaware of the location where her/his application is currently executed; it should believe that all the computations be local. This goal has been partially reached. Even though a Grid middleware has the purpose of hiding all the low level interactions taking place in a Grid system, in practice the user has still the responsibility to describe the structure of her/his application or, in the worst case, adapts it to the requirements of the underlying Grid. From the *ShareGrid*

---

\*This work has been supported by TOP-IX and the Piedmont Region Agency under the Innovation Development Program.

[13] project experience, we have learned that one of the most important aspect to consider for making a Grid project successful, is the social aspect, that is inducing a potential Grid user to use the Grid infrastructure as an everyday part of her/his work.

With this in mind, the ShareGrid project provides to its users community a Web portal, the *ShareGrid Portal*, in order to ease the execution and the monitoring of a user application in the Grid system.

The ShareGrid Portal is a Web based Grid portal that provides to its users the access to Grid services and resources through the ShareGrid middleware. Due to the nature of the current ShareGrid middleware, its primary focus is toward *Bag-of-Tasks* (BoT) [17] applications. A BoT application is a kind of parallel application whose tasks are independent from each other. Though this is one of the simplest kind of programming model for the Grid computing, there are several real world applications that adopt it. BoT applications include, but are not limited to, *parameter sweep applications* [16], which are applications structured as a set of multiple experiments each of which is executed with a distinct set of parameters. Parameter sweep applications can be viewed as a simple means of exploring the behaviour of a complex system through a series of parametric experiments. Monte-Carlo and discrete-event simulation are a typical example of parameter sweep applications. The wide applicability of this type of application contributed to the spread of the Grid computing in the scientific community [12]. In fact it is extensively used for carry on several experiments in many scientific areas such as: extra-terrestrial intelligence [6], protein folding [3], high-energy physics [4], just to name a few.

The remainder of this section is organized as follows. In §1.2 we provide a short list of the most important related works. In §1.3 we explain how the job submission works, both with and without the portal, and what benefits the ShareGrid Portal brings. In §1.4 we give an overview of the ShareGrid Portal architecture. Finally, in §1.4 we present future works.

## 1.2 Related Works

In this section we present some of similar and consolidated projects that are well-known to Grid community. The *GridSphere* project [25] is a portal framework which provides a *portlet*-based Web portal [11]; it supports various middlewares, like the *Globus toolkit* [20], *Unicore* [27], and *gLite* [24], through portlet components, called *Grid-Portlets*. The ShareGrid Portal is a portal framework too. The main difference is the mechanism used for supporting a new Grid middleware. While the GridSphere approach makes use of portlet components for extending the range of supported middlewares, the ShareGrid Portal extension

mechanism consists in a series of *Plain Old Java Object* (POJO) [22] interfaces, defining the high-level behaviour of a middleware, which are deployed in the portal by means of simple Java libraries (JARs). Another difference is that GridSphere delegates each middleware portlet for providing its user interface, while the ShareGrid Portal provides a uniform view independent by the underlying middleware.

The *P-GRADE Portal* [23] is a workflow-oriented Grid portal; it is built upon GridSphere, for the Web interface, *JavaGAT* [29], for interacting with the middleware, and *Condor DAGMan* [1] for managing a workflow application. The main difference with the ShareGrid Portal is the nature of the supported Grid applications; the P-GRADE Portal is oriented to workflow applications, with some extensions for parameter sweep applications; the ShareGrid Portal currently supports BoT applications, which include the family of parameter sweep applications but are more limited than workflow applications. Another difference is the type of the supported middleware. The P-GRADE Portal is a Globus-based, multi-Grid collaborative portal; it is able to connect to different Globus-based Grid systems and let their user communities to migrate applications between Grids. The ShareGrid Portal can be considered a multi-Grid portal as well; however, it is not limited to the Globus middleware, since it relies on a set of interfaces that abstract from the underlying middleware; regarding the multi-Grid collaboration aspect, it is relied upon the underlying middleware.

## 1.3 Job Submission

In this section we provide an overview of how the job submission works and the main motivations that brought us to develop a Web portal.

In order to submit an application to a Grid system, a user has to “prepare” a job that describes the structure and the behaviour of that application. In *OurGrid* version 3 [17], the middleware that the ShareGrid infrastructure currently adopts, the submission of a job is done through the *mygrid* program [18]. This is a Linux console application, installed on the user machine, that acts as a Grid local scheduler: it accepts in input a job file (i.e., the description of a user application) from the user and assigns it one or more computational resources according to a preconfigured scheduling policy. The file representing the user job is a text file following the *Job Description File* (JDF) format. Basically, a JDF file is a text file containing a collection of task specifications, each of which includes an optional “init” section (for the stage-in phase), a “remote” section (for the remote execution phase) and a “final” section (for the stage-out phase), as shown in Fig. 1. From the user point of view, the job submission phase requires three steps: (1) the user creates a job by writing a JDF file that provides the structure and the behaviour of the application (s)he wants to run, (2) the

JDF file is submitted to the mygrid program, which transparently takes care of mapping the related job on one or more computational resources, and (3) the user manually and periodically polls the mygrid program for checking the job execution status. This workflow has some weak points:

```

job:
  label: MyJob
  requirements: mem == 100MB
...
task:
  init:
    put MyLocal.in MyRemote.in
    ...
  remote: MyCommand -i MyRemote.in -o MyRemote.out
  final:
    get MyRemote.out MyLocal.out
    ...

```

**Figure 1. ShareGrid Portal – Example of a JDF file.**

- The mygrid program actually only runs on the Linux operating system; maybe, this is the major weakness since the computer world is not Linux-centric.
- The mygrid program maintains its state in the volatile memory of the client machine and so it must remain running for the entire duration of the job execution. For this reason, the user must keep her/his machine powered on until all of her/his jobs in execution are done. This point represents another important weakness causing both economical and ecological implications for the useless power consumption. In fact, the price of the energy continuously gets higher and higher due to the strong dependence to oil and to the current oil crisis caused by an unbalance between the demand and the offer. Hence, a possible economical consequence for a user is the increasing of her/his energy costs. For what concerning ecological implications, it is now a fact that the excessive energy consumption contributes to the global warming. Users that are sensible to this subject might be disappointed for seeing that, in some sense, they are contributing to the overheating of the planet Earth.
- The mygrid program is not free in resource occupation. We have empirically observed that when the mygrid program is in an idle state, the memory consumption is about of 30MB (the CPU, on the other hand, is nearly unutilized). For this reason, the user might consider the installation and the use of the mygrid program as something of too intrusive, especially when (s)he does not own a powerful machine.
- The user must repeatedly poll the mygrid application for knowing the execution status of a given job (e.g., either completed, failed or still running).
- The JDF syntax, though simple, is error-prone, especially for the beginner user; furthermore, some kind of error (e.g. a misspelled file name) might only be thrown near the end of the execution, making the entire computation useless.

The above issues were enough to motivate the realization of the ShareGrid Portal. The first benefit the Web portal brings is the operating system independence, freeing the user from having installed on her/his own machine the Linux operating system. A Web application has also the advantage that does not require any software installation, since almost all modern operating system distributions ship with a Web browser and the utilization of the resources of the user machine is rather limited. Because the Web portal does not store any job submission state on the user machine, the user can submit her/his jobs from any machine connected to the Internet; furthermore, there is no more need to keep the machine powered on, waiting for jobs completion, since all the informations about job submissions are kept in the ShareGrid infrastructure. In order to avoid the user to manually and periodically poll for monitoring the status of the job execution, the ShareGrid Portal provides an active notification system. When the execution status of a job changes (e.g. from running to finished), the portal sends a notification (actually an email) to the user.

The submission of a job has been simplified thanks to the presence of several user-friendly web interfaces. These interfaces divides in two main groups according to the way job informations are fed: (1) job file upload and (2) manually job insertion. The fastest way to submit a job through the portal is by using the “import file” interface, depicted in Fig. 2. This interface allows the user to directly upload a JDF job file, that is a file that contains informations about the structure and the behaviour of the user application. Along with the job file, the user can upload many input files as needed that will be transfered on worker machines during the stage-in phase. This type of job submission is targeted to expert users that do not want to go through the additional steps that are inevitably introduced by the others more user-friendly interfaces. The other type of job submission interface is the manual job insertion view. In this interface, the user can choose to insert her/his job between two views: (1) a generic simple interface and (2) an ad-hoc interface for parameter sweep applications. In the generic simple interface, shown in Fig. 3, the user can create a job for a generic BoT application. For each task, the user must explicitly specify the executable command line (i.e., the executable name and its argument that will be executed on the worker machine), can optionally upload input files (in-

(Text marked with \* is mandatory)

Import File Manual Insertion

Job File \*  Browse... Upload

Input Files  Browse... Upload

Submit Preview Export Cancel

**Figure 2. ShareGrid Portal – Job submission through job file import.**

cluded the executable command if not already present in the worker machines) and possibly specify one or more output file names. This last two informations are used during the job stage-in and stage-out phase, respectively. The other

(Text marked with \* is mandatory)

Import File Manual Insertion

Job Name

Job Requirements

Simple View  Parameter Sweep View

Input Files  Browse... Upload

Tasks Insertion Executable Command Line \*  Add Task

Output Files  Add

Submit Preview Export Cancel

**Figure 3. ShareGrid Portal – Job submission through the simple view.**

way for manually submit a job is the parameter sweep view, shown in Fig. 4, an ad-hoc interface targeted for parameter study applications. This kind of application differs from a generic BoT application for the executable command: possibly different for each task in a BoT application and unique in a parameter sweep application. In fact, in this interface, the user is asked to specify a unique executable command, the list of parameters to study (one line for each experiment) and zero or more input, output and shared files. In order to speed up the insertion of informations, this view provides many useful shortcuts; the ones that are worth noting are:

- The user can choose to manually insert the executable command name, in the case it is already installed on worker machines, or to upload the corresponding executable command file.

- The arguments for the executable command and the output file names can be either uploaded via a text file or inserted by hand.

- Parameters can be studied as a function of input files by combining each parameter line to every input files.

(Text marked with \* is mandatory)

Import File Manual Insertion

Job Name

Job Requirements

Simple View  Parameter Sweep View

The executable is already installed on the remote machines.

Executable \*  Browse... Upload

Import the arguments for the executable command from a text file.

Executable Arguments

Tasks Insertion

Specify input files as executable command arguments.  
 Combine each argument to every input files.

Input Files  Browse... Upload

Specify additional files as executable command arguments.

Additional Files  Browse... Upload

Specify output files as executable command arguments.  
 Associate to each task only one output file (order matters).  
 Import the output files from a text file.

Output Files  Add

Submit Preview Export Cancel

**Figure 4. ShareGrid Portal – Job submission through the parameter sweep view.**

Each of the above interfaces allows the user to preview the job before submitting it, for discovering possible syntax errors, and to export the job to a text file (actually, only to a JDF file); the exported job can be successively uploaded in the import file view, in order to let the user to minimize repetitive tasks for similar jobs. After a job has been submitted, the user is freed from any other task and can decide, for example, to submit another job or even to shutdown her/his machine. It is the responsibility of the ShareGrid Portal to instruct the underlying Grid middleware for staging in the input files, remotely executing the specified command and, finally, for staging out the output files. In particular, staged

out files are stored in the ShareGrid repository, that is an area accessible only to the user who submitted the job. Once the execution status of a job is changed (e.g. from running to finished or to failed), the ShareGrid Portal sends a notification to the user. For instance, when a job has been successfully completed, the notification includes a link to the ShareGrid repository where the output files has been stored.

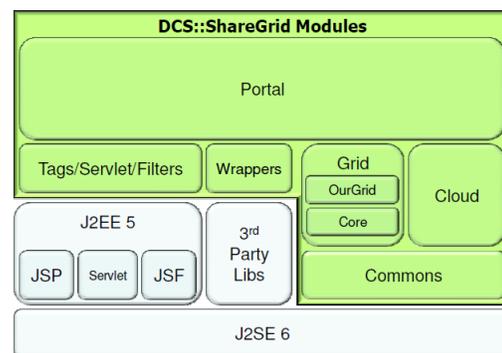
## 1.4 Architecture

In this section we provide a high-level description of the ShareGrid Portal architecture. The ShareGrid Portal is both a Web portal application and a portal framework.

As a Web portal application, it provides internationalization and localization support, user account management, data persistence abstraction, graphical appearance customization, and a set of core functionalities for the creation, deletion, updating and querying of user and job informations. The access control to the portal is based upon the *Role Based Access Control* (RBAC) model [19]. As pointed out in [28], a role is a semantic construct around which access control policies are formulated; users are assigned to specific roles and, in this way, they acquire the permissions associated with their roles. Roles are closely related to the concept of groups but the main difference is that a role brings together a set of users on one side and a set of groups on the other, while a group is typically defined as a set of users. Our RBAC model consists of a hierarchical user role model where each role is assigned to one or more access permissions defined upon a hierarchical page access control policy. For instance, the administrator user is allowed to access to any page whereas the anonymous user can access only to a limited set a pages (e.g., to the user registration page). Actually, we have defined three roles: (1) Anonymous, for users not logged into the portal, (2) Standard, for users that are allowed to submit a job to the ShareGrid infrastructure, and (3) Administrator, for standard users with additional site administration privileges.

As a portal framework, the ShareGrid Portal provides to the developer a set of independent and reusable components. In Fig. 5 is depicted an high-level view of the ShareGrid Portal infrastructure. All the modules relies on the *Sun Java Platform Standard Edition 6* [8]. The “Commons” module offers shared and commonly used functionalities, like string manipulation, format conversions, I/O and network utilities, and so on. This component is used by almost any other ShareGrid modules. The “Grid” module aims to provide an abstraction layer from any Grid middleware; the ShareGrid portal uses this component for keeping it independent by the underlying middleware used. This module is divided into two parts: the “Core” sub-module defines the interfaces and the implementations that are middleware independent, whereas the OurGrid sub-module is the im-

plementation for the OurGrid middleware. The role of the “Portal” module is two fold: (1) it provides a set of interfaces, classes and tag libraries to act as a Web application framework for developing *Sun Java EE* [7] Web applications, and (2) it realizes the Web interfaces for using it as a Grid portal. A Grid portal derived from the “Portal” module, included the ShareGrid Portal, consists of at least a set of presentation pages, including static *HTML* [26], *Sun JavaServer Pages* [10] and *JavaServer Faces* [9], along with the associated backing beans, for implementing the presentation logic; in addition, it is possible to define classes for realizing the business logic and overriding existing classes for redefining, for instance, the data persistence layer or the page life cycle.



**Figure 5. ShareGrid Portal – High level view of the architecture.**

## 1.5 Conclusions and Future Work

The ShareGrid Portal is a rather young project started in the middle of 2007. Nevertheless, it is able to provide a simple but effective way to submit jobs to the ShareGrid infrastructure avoiding to force its users to adapt their desktop environment to the requirements of the underlying middleware. Obviously, from the point of view of an expert user, the time taken for submitting a job to the Grid middleware with the portal will never be comparable with the one spent directly using the console application. In fact, this is a trade-off that almost all Web applications have to accept with respect to the desktop-based counterparts. However, we think the benefits brought by a Web portal might make the adoption of the ShareGrid infrastructure more attractive.

Being a young project it is in continue evolution. On-going projects include the redesign of some views, in order to make the job submission even faster, and the development of application oriented Web interfaces, that is interfaces specifically targeted to an application domain, like distributed rendering. Future extensions include the support

for others Grid middlewares, the implementation of a Web Services layer and the possibility to export a job to different formats, like the Job Submission Description Language (JSDL) [14] format.

## References

- [1] The Condor DAGMan. <http://www.cs.wisc.edu/condor/dagman/>.
- [2] European Organization for Nuclear Research (CERN). <http://www.cern.ch/>.
- [3] The Folding@home project. <http://folding.stanford.edu/>.
- [4] The LCG project. <http://lcg.web.cern.ch/LCG/>.
- [5] The LHC project. <http://lhc.web.cern.ch/lhc/>.
- [6] The SETI@home project. <http://setiathome.berkeley.edu/>.
- [7] The Sun Java Platform Enterprise Edition 5. <http://java.sun.com/javasee>.
- [8] The Sun Java Platform Standard Edition 6. <http://java.sun.com/javase>.
- [9] The Sun JavaServer Faces 1.2. <http://java.sun.com/javase/ava-server/faces/>.
- [10] The Sun JavaServer Pages 2.1. <http://java.sun.com/products/jsp/index.jsp>.
- [11] A. Abdelnur and S. Hepper. JSR 168: Java Portlet specification version 1.0. Technical report, Sun, 2003. <http://www.jcp.org/en/jsr/detail?id=168>.
- [12] D. Abramson, J. Giddy, and L. Kotler. High performance parametric modeling with Nimrod/G: Killer application for the global Grid? In *IPDPS '00: Proceedings of the 14th International Symposium on Parallel and Distributed Processing*, pages 520–528, Cancun, Mexico, May 2000. IEEE Computer Society.
- [13] C. Anglano, M. Canonico, M. Guazzone, M. Botta, S. Rabbellino, S. Arena, and G. Girardi. Peer-to-peer desktop Grids in the real world: The ShareGrid project. In T. Priol, L. Lefevre, , and R. Buyya, editors, *CCGRID'08: Proceedings of the 2008 Eighth IEEE International Symposium on Cluster Computing and the Grid (CCGRID)*, pages 609–614, Lyon, France, May 2008. IEEE Computer Society.
- [14] A. Anjomshoaa, F. Brisard, M. Drescher, D. Fellows, A. Ly, S. McGough, D. Pulsipher, and A. Savva. GFD.136: Job Submission Description Language (JSDL) specification, version 1.0 (first errata update) [obsoletes GFD.56]. Technical report, Open Grid Forum (OGF), July 2008. <http://www.ggf.org/documents/GFD.136.pdf>.
- [15] M. Baker, R. Buyya, and D. C. Hyde. Cluster computing: A high-performance contender. *IEEE Computer*, 32(7):79–80, 1999.
- [16] H. Casanova, G. Obertelli, F. Berman, and R. Wolski. The AppLeS parameter sweep template: User-level middleware for the Grid. *Scientific Programming*, 8(3):111–126, 2000.
- [17] W. Cirne, D. Paranhos, L. Costa, E. Santos-Neto, F. Brasileiro, J. Sauvé, F. A. B. Silva, C. O. Barros, and C. Silveira. Running Bag-of-Tasks applications on computational Grids: The MyGrid approach. *International Conference on Parallel Processing (ICPP)*, page 407, October 2003.
- [18] L. B. Costa, L. Feitosa, E. Araújo, G. Mendes, R. Coelho, W. Cirne, and D. Fireman. MyGrid: A complete solution for running Bag-of-Tasks applications. In *SBRC 2004: 22nd Brazilian Symposium on Computer Networks (SBRC)*, Gramado, RS, Brazil, May 2004.
- [19] D. Ferraiolo and R. Kuhn. Role-Based Access Control. In *15th NIST-NCSC National Computer Security Conference*, pages 554–563, 1992.
- [20] I. Foster and C. Kesselman. Globus: A metacomputing infrastructure toolkit. *The International Journal of Super-computer Applications and High Performance Computing*, 11(2):115–128, Summer 1997.
- [21] I. Foster and C. Kesselman. *The Grid: Blueprint for e New Computing Infrastructure*. Morgan Kaufmann, 1998.
- [22] M. Fowler. POJO: An acronym for Plain Old Java Object. <http://www.martinfowler.com/bliki/POJO.html>.
- [23] P. Kacsuk and G. Sipos. Multi-Grid, multi-user workflows in the P-GRADE grid portal. *Journal of Grid Computing*, 3(3):221–238, September 2005.
- [24] E. Laure, F. Hemmer, A. Aimar, M. Barroso, P. Buncic, A. D. Meglio, L. Guy, P. Kunszt, S. Beco, F. Pacini, F. Prelz, M. Sgaravatto, A. Edlund, O. Mulmo, D. Groep, S. Fisher, and M. Livny. Middleware for the next generation Grid infrastructure. In A. Aimar, J. Harvey, and N. Knoors, editors, *CHEP 2004: Computing in High Energy Physics and Nuclear Physics (CHEP)*, page 826, Interlaken, Switzerland, Sep. 2004.
- [25] J. Novotny, M. Russell, and O. Wehrens. GridSphere: A portal framework for building collaborations: Research articles. *Concurrency and Computation: Practice & Experience*, 16(5):503–513, 2004.
- [26] D. Raggett, A. L. Hors, and I. Jacobs. Html 4.01 specification. Recommendation, World Wide Web Consortium (W3C), December 1999.
- [27] M. Romberg. The UNICORE Grid infrastructure. In *Proceedings of 1st Worldwide SGI Users' Conference*, pages 144–153, Krakow, Poland, 2000.
- [28] R. S. Sandhu, E. J. Coyne, H. L. Feinstein, and C. E. Youman. Role-Based Access Control models. *IEEE Computer*, 29(2):38–47, 1996.
- [29] R. V. van Nieuwpoort, J. Maassen, R. Hofman, T. Kielmann, and H. E. Bal. Ibis: An efficient Java-based Grid programming environment. In *Joint ACM Java Grande - ISCOPE 2002 Conference*, pages 18–27, Seattle, Washington, USA, November 2002.