

A Grid Portal for an Undergraduate Parallel

Programming Course

Juan Touriño, *Member, IEEE*, María J. Martín, Jacobo Tarrío, Manuel Arenaz

Computer Architecture Group

Department of Electronics and Systems, University of A Coruña, Spain

Corresponding author: Juan Touriño, Department of Electronics and Systems,
University of A Coruña, Campus de Elviña, s/n, 15071 A Coruña, Spain (e-mail:
juan@udc.es)

Abstract

This paper describes an experience of designing and implementing a portal to support transparent remote access to supercomputing facilities to students enrolled in an undergraduate parallel programming course. As these facilities are heterogeneous, are located at different sites, and belong to different institutions, grid computing technologies have been used to overcome these issues. The result is a grid portal based on a modular and easily extensible software architecture that provides a uniform and user-friendly interface for students to work on their programming lab assignments.

Index Terms: Grid computing, grid portal, supercomputing, parallel programming

1 Introduction

Since 1998, the Department of Electronics and Systems (DES) at the University of A Coruña (region of Galicia, Spain) has offered a one-semester elective course in parallel programming for undergraduate fifth-year Computer Science (CS) students. The inclusion of parallel processing in an undergraduate degree has been widely justified since the eighties [1] since parallel computers have become much easier to use and much more widely available. It continued to be justified in the nineties [2]-[4], and nowadays many universities offer specific courses to provide CS undergraduates with real parallel processing experience [5] as an essential component to teaching parallel programming concepts effectively. Three additional reasons have favored the introduction of a parallel programming course in the CS curriculum at the University of A Coruña: (a) there is an active research group on parallel processing at DES; (b) there are powerful supercomputers available at the Supercomputing Center of Galicia (CESGA) [6] in Santiago de Compostela (70 kilometers away from DES), mainly used for scientific computing; (c) there is growing interest of small and medium enterprises of Galicia in cluster architectures [7] because they deliver outstanding parallel performance at a competitive cost. CS graduates skilled in parallel processing techniques are, therefore, requested by both academic/scientific and enterprise environments.

Problems that arise in teaching a course in parallel programming include the difficulty in having access to supercomputing resources and the time it takes to teach their effective use, as stated in [8]. Wilkinson and Allen discourage the use of resources of supercomputing centers [9] because the technicalities of supercomputers may be too much for an undergraduate course. Moreover, their access must be controlled carefully by the center, and perhaps the centers are not equipped

to handle a large number of inexperienced students. Therefore, the choice is the use of a dedicated local network of workstations for parallel programming.

Since the course was intended to be highly practical, a hybrid approach was followed: use of local resources (both from DES and the School of Computer Science (SCS) at the University of A Coruña), mainly for training purposes, and access to remote supercomputers (available at CESGA) to assess the performance of parallel algorithms. This approach has several advantages: students are more motivated if they have access to real supercomputers; and they have the chance to implement the same algorithm using different architectures and parallel programming models, which enable students to make objective performance comparisons (execution times, speedups, complexity of algorithm development, etc.). Moreover, their curiosity for new machines stimulates self-learning to go deeply into specific parallel processing issues of the target supercomputers.

The computers used in the course are administered by different institutions: DES, SCS, and CESGA. As a minor drawback, tedious bureaucratic procedures are necessary to obtain accounts for the students (mainly, access to the supercomputing center). However, the major drawback is the impact on students in using a number of unrelated and geographically distributed systems. Thus, each student has several accounts on machines with their own local file system (which involves source code files being continuously transferred among the target machines), and with different access, use, and security policies (e.g., restricted IP access that makes it difficult to login to the machines from home). Once the student is logged on, he/she has to deal with different compilers, libraries to be linked, and specific program execution commands for each machine and each parallel programming model, as well as different job schedulers to submit and monitor parallel jobs.

All these issues outside the core course contents involve significant overhead time. These

problems were overcome in 2003 at the University of A Coruña by applying grid computing technologies [10] to facilitate the use of the diverse parallel processing platforms. Specifically, a grid portal was developed that acts as a single point of entry to the geographically distributed computers used in the course and as a high-level, user-friendly environment for the students to manage transparently their parallel jobs. No reported practical experiences on the application of a grid portal to support the teaching of a course were found, but a few projects focused on setting up a grid infrastructure for teaching purposes, such as ISILab [11] and ULabGrid [12]. ISILab is a distributed environment that allows students to carry out remote lab experiments with real electronic instruments and circuits via the web. The authors propose to map ISILab functionalities on a layered grid model: from the grid infrastructure composed of diverse resources geographically scattered (computers, storage systems, catalogues, networks, sensors, devices) to the user application layer that controls the experiment execution. But this grid-oriented approach is under development using proprietary protocols and APIs, which limit reusability for other projects. ULabGrid is an architecture based on standard grid technologies that enables educators to design collaborative, distant laboratories for students using a grid infrastructure. It allows students to run the software required for the lab from everywhere at any time. But this approach was not conceived to access supercomputing facilities, and it does not provide web access via a portal. A prototype of this architecture applied to a flight simulator of an aeronautics lab is still being developed.

The rest of the paper is organized as follows: Section 2 describes the contents of the course (both lectures and laboratory work) which provide the context for the design and development of the grid portal. Section 3 focuses on the portal: requirements, software architecture components and functionalities provided to ease the completion of the course assignments. Finally, conclusions are

drawn in Section 4.

2 Course content

The one-semester course consists of 60 hours (40 hours in-class instruction and 20 hours laboratory). It attracts 20 students per year on average. All the students in the fourth year take a course in advanced computer architecture which provides the architectural basis (including performance issues) for the parallel programming course. Thus, the students have a solid background in advanced topics, such as parallel computer architectures (shared memory, distributed memory), interconnection networks for parallel computers, cache coherence and memory consistency on multiprocessors, etc. Students are also skilled in C and Fortran 90 (taught in previous courses), which are the prerequisite languages for the course. Prior to the introduction of the parallel programming course in 1998, the advanced computer architecture course included a short parallel programming lab focused on the Parsys Supernode, a transputer-based machine (from 1990 to 1993), and PVM [13] message-passing programming on a network of workstations (1994-1997).

2.1 Course lectures

The syllabus for the parallel programming course currently being offered is on the following lines.

1. Introduction

- Course organization
- Review of parallel computer architectures

- Parallel programming models
- Case study: computing π in parallel using numerical integration

2. Message-passing programming: MPI

- MPI main features
- Point-to-point communications
- Collective communications
- Derived data types
- Communicators and virtual topologies
- Overview of new functionalities in MPI-2

Recommended textbook: [14]. Web material: [15]

3. Shared memory programming: OpenMP

- OpenMP main features
- Directives to exploit loop-level parallelism
- Directives for task-level parallelism
- Synchronization mechanisms
- Runtime library and environment variables

Recommended textbook: [16]. Web material: [17]

4. Data-parallel programming: HPF

- HPF main features
- Data distribution and alignment directives
- Directives and constructs to express data parallelism
- HPF intrinsic and library procedures
- Extrinsic procedures

Recommended textbook: [18]. Web material: [19]

5. Parallel programming techniques: design of parallel programs

- Performance metrics for parallel programs
- Data partitioning techniques
- Techniques to enhance locality: exploiting memory hierarchy
- Load balancing techniques
- Case studies: parallel sorting, matrix multiplication, direct and iterative methods for solving linear systems of equations, parallel FFT, parallel tree search

Recommended textbook: [20]. Web material: [21]

The course ends with Lecture 5 because its contents may be too abstract for students at the beginning of the course without having experienced practical parallel programming using the specific software taught in Lectures 2-4. Nevertheless, some topics of Lecture 5 are informally or intuitively introduced in the previous lectures, for instance: basic performance metrics such as speedup or efficiency, load balancing techniques through OpenMP scheduling clauses, or data partitioning

strategies using HPF directives. With this basis set this lecture consolidates and completes the learning of parallel programming techniques through practical case studies, thus these topics are more easily assimilated by students.

2.2 Course laboratory

Lab work consists of the following assignments.

1. Training exercises. The students are provided with the skeleton of several example codes of increasing difficulty to illustrate and experiment with the main concepts explained for each programming model.
2. Comparison of programming models. The students have to implement the same parallel algorithm (e.g., Gauss-Jordan elimination in the 2003 course) using MPI, OpenMP, and HPF in order to compare practically the programming features associated with each model.
3. Algorithm performance contest. A well-specified algorithm is proposed to be implemented in parallel (using MPI, since it provides more programming flexibility) on a particular super-computer. The goal of the student is to maximize performance by taking into account both programming and architectural issues. Some examples used in previous years are parallel sorting algorithms, parallel matrix transposition, data redistributions (e.g., block-to-cyclic), or parallel prefix/suffix operations. This contest stimulates creativity, in-depth study, as well as experience sharing, discussion, and criticism (the best works are presented in class).
4. Course miniproject. It is an individual work proposed by the student under the supervision of

Table 1: Supercomputing resources available for the students

Computer	Location	Programming	Short description
PC Network (sky.des.udc.es)	DES	MPI (MPICH)	10 PCs (PIII 800 MHz), Fast-Ethernet network
Self-made SCI cluster (muxia.des.udc.es)	DES	MPI (ScaMPI)	10 racked dual-processor nodes (PIV Xeon 1.8GHz), SCI network
Compaq Beowulf cluster (bw.cesga.es)	CESGA	MPI (MPICH-GM) HPF	16 racked single-processor nodes (PIII 1GHz), Myrinet 2000 network
SGI Origin 200 (silgar.cecafi.fi.udc.es)	SCS	OpenMP	ccNUMA multiprocessor (4 MIPS R10000 180 MHz)
HP Superdome (sd.cesga.es)	CESGA	OpenMP	cluster of 2 64-processor ccNUMA systems (Intel Itanium2 1.5 GHz)
Compaq HPC320 (sc.cesga.es)	CESGA	OpenMP/MPI hybrid	cluster of 8 4-processor SMP nodes (Alpha EV68 1 GHz)

the teacher who determines if it has enough technical merit. The goal of this miniproject is to put into practice as many theoretical contents as possible through the implementation of a parallel algorithm, application, or library using a particular programming model. Message-passing is usually the choice because students can install MPICH [15], a free implementation of MPI, on their own PCs for an initial development of their codes before moving them to a parallel computer. Students are encouraged to work on topics that have been studied in other courses (thus these topics are familiar to them) to enhance interdisciplinary education. The most popular fields selected for the miniproject are numerical computation, neural networks, genetic algorithms, computer graphics, and image processing.

The resources currently available for the lab assignments are shown in Table 1. The “Programming” column shows the parallel programming software used for each computer, which does not indicate the only way to program that computer. Local resources are mainly devoted to training exercises and to the development and debugging of the other programming assignments before being

tuned, executed, and evaluated on a particular CESGA supercomputer.

3 Grid portal for the course lab

The lab resources listed in Table 1 (which may vary each year depending on availability) are heterogeneous and are administered by different institutions. These resources have been integrated into a grid environment for the parallel programming lab. Grid computing [10] links different computing resources that belong to different institutions (usually sited in distant physical locations) and makes these resources available to users located at remote sites via public communication networks. Regarding this educational scope, the grid must provide the following basic functionalities.

- **Authentication:** students and computing resources must be authenticated to allow the former to connect to the latter securely.
- **Information service:** an information service must provide students with both static and dynamic information about the setup and the status of the hardware and software resources of the grid (e.g., resource characteristics and availability).
- **Data transfer:** students must be able to transfer securely and transparently their source codes and associated data files between the different computers of the grid.
- **Resource/job management:** a unified mechanism must provide for resource allocation, remote job submission, cancellation, monitoring, job input/output, etc.

These functionalities are provided by diverse software components which can be found integrated in software packages called grid toolkits. The Globus toolkit [22], which aims to use com-

modity software and technologies, is becoming the de facto standard for grid computing. It provides middleware services to support grid environments. These toolkits are designed to provide basic grid services through low-level command-line interfaces, and thus a significant amount of learning time and effort must be spent by their potential users. This situation would be a serious obstacle to the effective adoption of the grid for instruction support, since the goal is precisely to enable students to develop more easily their assignments on the lab resources. To alleviate this problem, some easy and intuitive interfaces, such as grid portals that take the complexity of the grid away from the user, have been developed. A grid portal is a web-based single point of entry to a grid and its implemented services. The Grid Portal Development Kit (GPDK) [23], a set of higher-level function libraries that provide a generic framework for the development of grid portals, was used to implement the portal. The GridPort toolkit [24] also provides access to grid services through a web interface by means of a set of Perl modules which wrap Globus. But the design of GridPort makes code maintainability and portal customization more difficult than the approach taken by GPDK. A recent promising successor to GPDK is the GridSphere portal framework [25] with a new architecture based on portlets. The goal of this framework is to enable portal developers to implement new portal interfaces and add new functionality more quickly.

3.1 Portal design and architecture

A customized grid portal for lab support has been designed using several standard software components organized as a vertical stack, where each layer uses the services provided by the lower layer (Fig. 1). In the upper layer, the user opens a connection to the secure Apache web

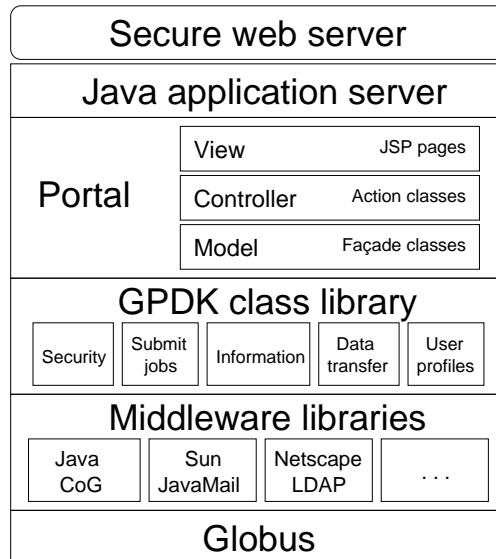


Figure 1: Software architecture of the grid portal

server and makes an HTTP request. This request is forwarded to the Tomcat application server, which is running the portal web application. Although GPKD provides a set of templates to generate ready-to-run portals, they have not been used since they implement an ad hoc architecture based on the use of program logic embedded in JSP pages, which presents many drawbacks. For example, customization of the portal look for a particular site would probably involve changing some lines of Java code, which is generally outside of what a website designer is willing to do. Conversely, adding new functionalities to the portal would require the portal developer to design its appearance at the same time. To avoid this issue, the Model-View-Controller (MVC) architecture [26] was used to provide a greater separation between business logic and presentation. The MVC architecture was used to design a new layer on top of GPKD (the Portal layer depicted in Fig. 1) and was implemented using the Java Struts framework. In the MVC architecture, the Model contains data and functionalities (e.g., file transfer, job submission...); the View is the part of the portal that interacts with the user (graphical display); and the Controller connects View and Model together and controls

the execution of the portal. The Model was implemented using customized, high-level methods that make calls to functions of the lower layer, the GSDK class library that provides access to grid services along with some portal-specific utilities (e.g., portal user profiles). GSDK uses the libraries of the next layer: the Java Commodity Grid (CoG) kit [27] (which provides Java counterparts of the Globus C APIs) and some commodity libraries not specifically created for grid computing, but used to access grid services (e.g., the Globus information service is based on the LDAP protocol). Finally, in the lowest layer lies Globus, which provides the basic grid services.

This layered approach allows to add new computing resources straightforwardly. When a new computer joins the grid, the associated information (e.g., number of processors or job managers installed) is automatically transferred from the Globus information service at the lowest layer of Fig. 1 to the upper layers, and the Portal View layer displays the interface to have access to the new resource. Moreover, the strong separation between layers provided by the MVC architecture allows the modification of the appearance of the portal (HTML code) without affecting the Java code that implements the functionalities and vice versa, and allows the inclusion of new functionalities easily.

3.2 Portal features and functionalities

The portal described in the previous subsection provides easy and transparent access to the available lab resources enumerated in Table 1. The basic functionalities of the portal are depicted as a navigation diagram in Fig. 2. It is a UML state diagram in which each portal page is represented by a state, and each link between pages is represented as a transition between states. The portal also includes links to useful course information: course schedule and news, and course material

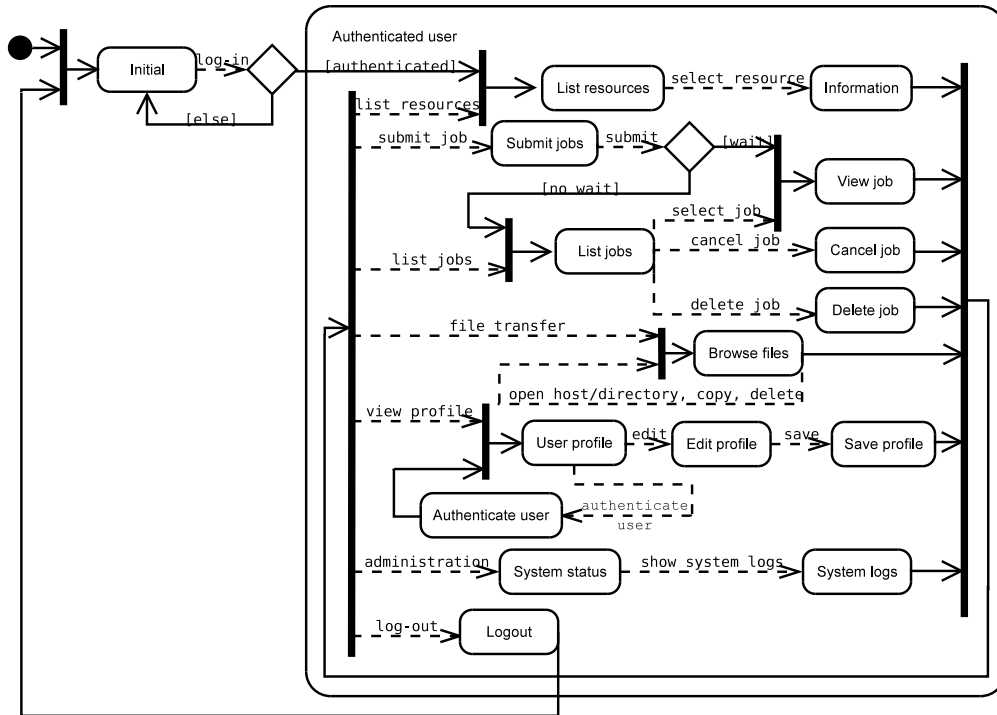


Figure 2: Navigation diagram of the grid portal

(local and external) to enhance self-study.

The portal home page is shown in Fig. 3. Students must be authenticated and authorized to access the grid resources. The DES system administrator issues and stores credentials (a certificate and a private key with a lifetime of one semester) for each student on a credential server. There are two alternative ways for students to log into the portal: using a username and password (with a maximum lifetime of one week for security reasons), obtained by running the MyProxy package [28] on the credential server; or using a proxy credential file that must be copied to the computer where the student will launch the browser to access the portal (usually his/her home computer). This file, also with a limited lifetime, is obtained by running the grid-proxy-init Globus command on the credential server.

Once the student has been authenticated in the portal, the user's home page is displayed (see Fig. 4). This page shows the available resources (which could be different depending on the user) and detailed resource-level information, such as available processors and memory, average system load, or description of available job queues. This information is obtained dynamically from the Globus information service. The student has access to different functionalities from this page, such as file transfer between computers of the grid, job submission, or user profile information (personal data, certificate details, user preferences, user history). For illustrative purposes, Fig. 5 shows the job submission page. It can be used to execute any program (OS commands, compilation commands, parallel codes, etc.) on the selected machine of the grid. In this example, an 8-processor MPI job is submitted to the muxia cluster. This computer provides two different job managers (information obtained dynamically from the grid and displayed in a menu): fork, for running interactive jobs, and PBS (Portable Batch System), which was selected to launch the parallel job. The link "List jobs" in the menu displays a complete list of the user's jobs (see Fig. 6): job name, execution node, current status (active or completed), submission date, and a link to the job output (if available). There is also a link to cancel a running job or to delete an entry from the list when the job is finished.

Additionally, the course teachers have access to several parameters about the portal activity ("System status" link): users currently connected, list of active and finished jobs, files transferred per user, login/logout records, and other activity logs. These parameters provide useful information on the individual and global use of the different computers of the grid. This information also turned out to be useful to detect and discourage incidents and bad practices in the use of supercomputers by students (e.g., massive submission of jobs, over-allocation of processors and memory, hung jobs

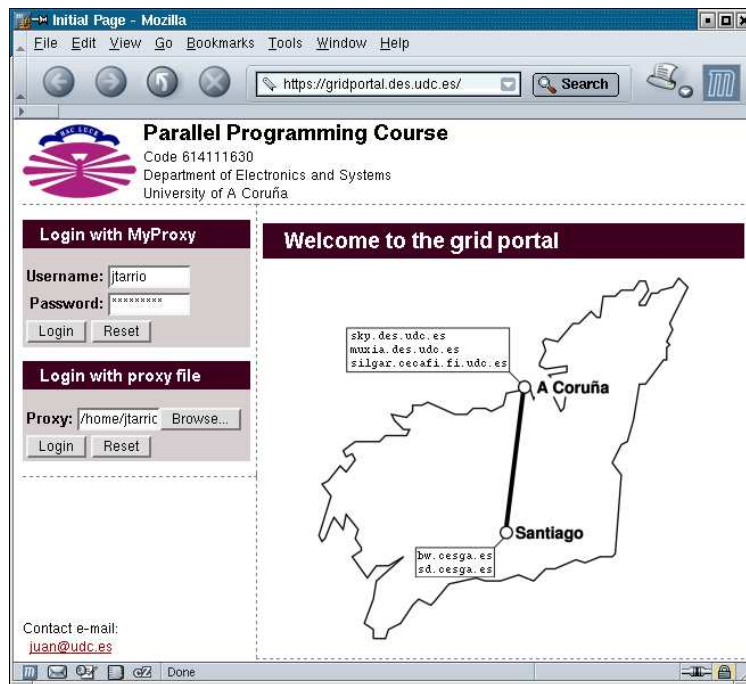


Figure 3: Portal home page with authentication

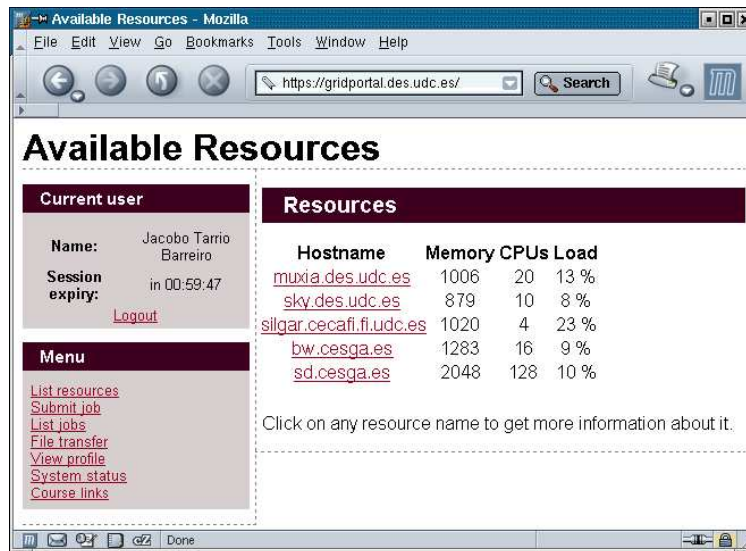


Figure 4: User's home page: resource list and menu of functionalities

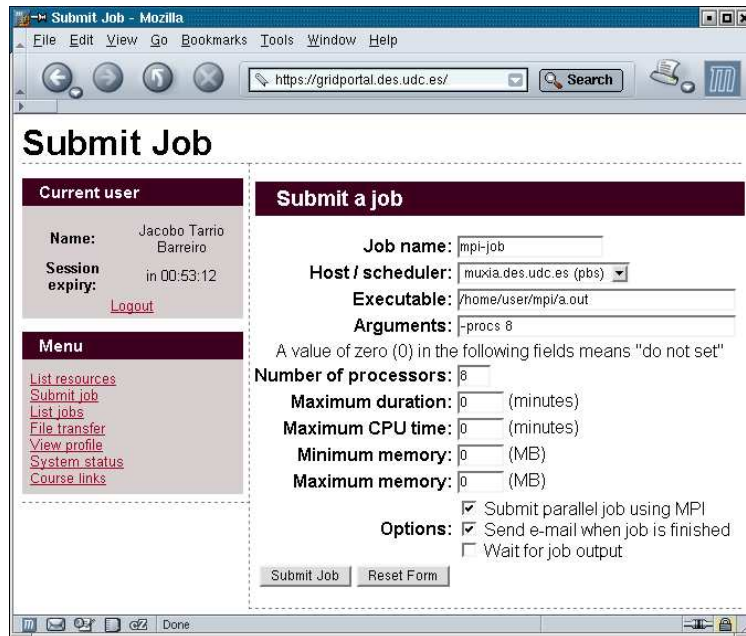


Figure 5: MPI job submission through the portal



Figure 6: List of user's jobs

as a result of careless job running).

Regarding the dynamics of the lab, students first develop and debug their assignments using local resources and submit their jobs to these resources in interactive mode through the portal. These jobs usually have small workloads and input data sets to check code correctness. Next, source codes are moved to CESGA supercomputers through the file transfer page. In the first contact with a real supercomputing environment, many students become impatient and complain because they are accustomed to obtain immediate results in the interactive execution on local resources. But now their jobs are submitted to queues through the job submission page and, depending on the load of the supercomputer (which can be checked from the resource portal page), some time is required to gain resource access since CESGA supercomputers are used by researchers from the three Galician universities and from CSIC, a national research institute. Students are advised to finish their assignments far enough in advance to allow for supercomputers that may be overloaded close to the assignment deadline. Nonetheless, students rapidly adapt to this new environment and experience the power of supercomputers by using more processors and larger workloads for their jobs. They also use specific profilers (usually linked to their codes) that study the application behavior and report performance statistics (e.g., CPU time, communication time, cache behavior) to guide performance tuning. At this stage, students become fully conscious that programming portability does not mean performance portability because performance results largely depend on the target machine. Students are provided in the lectures with basic information about available profiling tools or compiler options for each programming model, but they must experiment with them on their own. The effective use of these tools is an additional factor for the teacher to evaluate the assignments.

4 Conclusions

In this paper, a grid portal to support an undergraduate parallel programming course was described. Although grid technology is primarily intended to enable large-scale scientific research projects to better utilize and share distributed resources, it has been successfully applied in an educational environment. The portal provides an intuitive and homogeneous interface to the geographically distributed and dynamically changing supercomputing resources available for the lab assignments. Thus, the portal allowed students to focus on the course subjects instead of dealing with irrelevant and time-consuming technical issues, such as job submission or file transfer between computing resources, which helped to improve student performance.

The experience of the authors has shown that it is feasible for undergraduates to have access to supercomputing facilities, and that their assignments do not interfere with the research activities of the supercomputing center. This situation results because of the relatively small number of students enrolled in the course, and in general, the assignments are designed so that they are not very CPU and memory consuming (a maximum of a few CPU hours, reasonable enough for students to exploit the power of the supercomputer). Moreover, CESGA enthusiastically supports the course as a medium to promote and extend the use of parallel computing technologies in Galicia.

The introduction of the portal was valued very highly by the students, who have even provided feedback on the portal design and functionalities. They highlighted the 24-hour access to supercomputing facilities from home requiring only a browser, unlike other courses where hardware or software resources can only be accessed from on-campus machines. This fact, together with the motivation of programming production supercomputers, caused the students to spend many extra

hours at home on the assignments, as revealed by the statistics provided by the portal. The format of the lab assignments described in Section 2.2 also stimulated self-study (mainly the course miniprojects) which, in some cases, became research. In fact, this course turned out to be a good way to recruit Ph.D. students to work in the area of parallel processing.

Although the portal was developed for CS undergraduates, it could be used at the graduate level not only for CS students but also for students from diverse disciplines of science and engineering (such as mathematics, physics, chemistry, biology, and civil, mechanical, or aerospace engineering), enrolled in a computational science/engineering program, who need parallel programming as an essential tool to solve large-scale problems in their fields. The portal also opens new possibilities for teaching parallel programming in the context of international Master's and Ph.D. programs as a means to enable universities from different countries share their supercomputing resources for teaching purposes. It could also provide a remote parallel programming lab to students in countries that lack supercomputing infrastructures. The traditional educational links between Spain and Latin America offer a potential scenario for these purposes. Ongoing initiatives currently under development such as Gridcole [29], a collaborative learning environment that uses grid services technology, would be potentially appropriate in this context to provide integration of supercomputing capabilities or other specific resources at different locations in a wide area.

The development of the portal required a significant learning curve to acquire background in the technologies shown in Fig. 1. But grid computing is still evolving, and it is expected that grid tools will be easier to set up and use so that they are more accessible to a wide audience not familiar with grid technologies but who might wish to build a grid portal for teaching purposes. Once built, a key issue is to obtain organizational support to manage the portal. Thus, currently the departmen-

tal system administrator performs basic portal management tasks, such as portal accounting and authentication (e.g., student credentials), portal availability (maintenance of web, application and credential servers), or incident reporting on the grid infrastructure (e.g., unavailability of supercomputing resources). The URL for the portal home page is <http://gridportal.des.udc.es>.

Acknowledgments

This work was supported by a Teaching Innovation Grant (UDC-TIC03-057) from the University of A Coruña, by the Galician Government (Xunta de Galicia, Projects PGIDIT02TIC00103CT and PGIDIT04TIC105004PR), and by the CrossGrid European Project (IST-2001-32243). We gratefully thank CESGA for providing access to its supercomputing facilities. Also thanks to the Editor-in-Chief and three anonymous referees for their helpful comments.

References

- [1] Butler, R.M., Eggen, R.E., and Wallace, S.R., "Introducing parallel processing at the undergraduate level," *Nineteenth ACM Technical Symposium on Computer Science Education*, Atlanta, GA, pp. 63-67, Feb. 1988.
- [2] Hintz, T., "Introducing undergraduates to parallel processing," *IEEE Transactions on Education*, vol. 36, no. 1, pp. 210-213, Feb. 1993.
- [3] Toll, W.E., "Decision points in the introduction of parallel processing into the undergraduate curriculum," *Twenty-sixth ACM Technical Symposium on Computer Science Education*, Nashville,

TN, pp. 136-140, Mar. 1995.

- [4] Nevison, C.H., "Parallel computing in the undergraduate curriculum," *IEEE Computer*, vol. 28, no. 12, pp. 51-56, Dec. 1995.
- [5] Adams, J., Nevison, C., and Schaller, N.C., "Parallel computing to start the millennium," *Thirty-first ACM Technical Symposium on Computer Science Education*, Austin, TX, pp. 65-69, Mar. 2000.
- [6] Supercomputing Center of Galicia, <http://www.cesga.es> [Last visited: Nov. 2004].
- [7] IEEE Task Force on Cluster Computing, <http://www.ieeetfcc.org> [Last visited: Nov. 2004].
- [8] Youssefi, J.A., and Zemoudeh, K., "A course in parallel processing," *IEEE Transactions on Education*, vol. 40, no. 1, pp. 36-40, Feb. 1997.
- [9] Wilkinson, B., and Allen, M., "A state-wide senior parallel programming course," *IEEE Transactions on Education*, vol. 42, no. 3, pp. 167-173, Aug. 1999.
- [10] Foster, I., and Kesselman, C. (editors), *The Grid: Blueprint for a New Computing Infrastructure*. San Francisco, CA: Morgan Kaufmann Publishers, 1998.
- [11] Bagnasco, A., and Scapolla, A.M., "A grid of remote laboratory for teaching electronics," *Second International LeGE-WG Workshop on e-Learning and Grid Technologies*, Paris, France, Mar. 2003.
- [12] Ardaiz, O., Artigas, P., D'í az de Cerio, L., Freitag, F., Gallardo, A., Messeguer, R., Navarro, L., Royo, D., and Sanjeevan, K., "ULabGrid, an infrastructure to develop distant laboratories for

undergrad students over a grid," *First European Across Grids Conference*, Santiago de Compostela, Spain, Feb. 2003, Lecture Notes in Computer Science, vol. 2970, pp. 265-272, 2004.

[13] Geist, A., Beguelin, A., Dongarra, J., Jiang, W., Manchek, R., and Sunderam, V., *PVM: Parallel Virtual Machine. A Users' Guide and Tutorial for Networked Parallel Computing*. Cambridge, MA: The MIT Press, 1994.

[14] Pacheco, P.S., *Parallel Programming with MPI*. San Francisco, CA: Morgan Kaufmann Publishers, 1997.

[15] The Message Passing Interface (MPI) Standard, <http://www-unix.mcs.anl.gov/mpi> [Last visited: Nov. 2004].

[16] Chandra, R., Dagum, L., Kohr, D., Maydan, D., McDonald, J., and Menon, R., *Parallel Programming in OpenMP*. San Francisco, CA: Morgan Kaufmann Publishers, 2001.

[17] OpenMP Architecture Review Board, <http://www.openmp.org> [Last visited: Nov. 2004].

[18] Koelbel, C.H., Loveman, D.B., Schreiber, R.S., Steele Jr., G.L., and Zosel, M.E., *The High Performance Fortran Handbook*. Cambridge, MA: The MIT Press, 1994.

[19] The High Performance Fortran Home Page, <http://dacnet.rice.edu/Depts/CRPC/HPFF> [Last visited: Nov. 2004].

[20] Wilkinson, B., and Allen, M., *Parallel Programming Techniques and Applications Using Networked Workstations and Parallel Computers*. Upper Saddle River, NJ: Prentice-Hall, 1999.

- [21] Foster, I., *Designing and Building Parallel Programs*. Boston, MA: Addison-Wesley, 1995, <http://www-unix.mcs.anl.gov/dbpp/> (book and resources on-line) [Last visited: Nov. 2004].
- [22] Foster, I., and Kesselman, C., "Globus: A metacomputing infrastructure toolkit," *International Journal of Supercomputer Applications*, vol. 11, no. 2, pp. 115-128, Summer 1997, <http://www.globus.org> [Last visited: Nov. 2004].
- [23] Novotny, J., "The grid portal development kit," *Concurrency and Computation: Practice and Experience*, vol. 14, nos. 13-15, pp. 1129-1144, Nov./Dec. 2002, <http://www.doesciencegrid.org/projects/GPDK> [Last visited: Nov. 2004].
- [24] Thomas, M., Mock, S., Dahan, M., Mueller, K., Sutton, D., and Boisseau, J.R., "The Grid-Port toolkit: A system for building grid portals," *Tenth IEEE Symposium on High Performance Distributed Computing*, San Francisco, CA, pp. 216-227, Aug. 2001, <https://gridport.npaci.edu> [Last visited: Nov. 2004].
- [25] Novotny, J., Russell, M., and Wehrens, O., "GridSphere: A portal framework for building collaborations," *First International Workshop on Middleware for Grid Computing, ACM/IFIP/USENIX International Middleware Conference*, Rio de Janeiro, Brazil, pp. 178-185, Jun. 2003, <http://www.gridsphere.org> [Last visited: Nov. 2004].
- [26] Singh, I., Stearns, B., and Johnson, M., *Designing Enterprise Applications with the J2EE Platform*. Boston, MA: Addison-Wesley, 2002.
- [27] von Laszewski, G., Foster, I., Gawor, J., and Lane, P., "A Java commodity grid kit," *Concurrency and Computation: Practice and Experience*, vol. 13, nos. 8-9, pp. 645-662, Jul./Aug. 2001.

- [28] Novotny, J., Tuecke, S., and Welch, V., "An online credential repository for the grid: MyProxy," *Tenth IEEE International Symposium on High Performance Distributed Computing*, San Francisco, CA, pp. 104-111, Aug. 2001.
- [29] Bote-Lorenzo, M.L., Hernández-Leo, D., Dimitriadis, Y.A., Asensio-Pérez, J.I., Gómez-Sánchez, E., Vega-Gorgojo, G., and Vaquero-González, L.M., "Towards reusability and tailorability in collaborative learning systems using IMS-LD and grid services," *Advanced Technology for Learning*, vol. 1, no. 3, pp. 129-138, 2004.