

Virtual Lab for fMRI: Bridging the Usability Gap

Silvia D. Olabarriaga, Piter T. de Boer, Adam Belloum, Ketan Maheshwari
Informatics Institute, University of Amsterdam
{silvia,adam,ptdeboer,kmaheshw}@science.uva.nl

Jeroen G. Snel, Aart J. Nederveen
Academic Medical Center, University of Amsterdam
{j.g.snel,a.j.nederveen}@amc.uva.nl

Maurice Bouwhuis
SARA Computing and Networking Services
bouwhuis@sara.nl

Abstract

Grid technology can offer a powerful infrastructure for a broad spectrum of (scientific) application areas, but the uptake of grids by “real” applications has been slow. Several aspects contribute to this scenario, among them a large gap between the communities that develop and use the technology. While grid developments focus mostly on functionality, usability issues are also very important for the end-users. In the VL-e project, one of the goals is to bridge this gap, bringing (existing) grid technology to end-users in various application domains. This paper describes the current attempts to provide an IT infrastructure to facilitate neuroimaging research, in particular functional Magnetic Resonance Imaging (fMRI), which enables the study of brain function. The infrastructure offers resources for fMRI data acquisition, data storage, computing and services. All (local and grid) resources are accessed directly by the users from a virtual desktop implemented by a Virtual Resource Browser (VBrowser).

1. Introduction

Grid technology can be used to build IT infrastructures for performing complex tasks, being potentially powerful in a broad spectrum of (scientific) application areas. The uptake of grids by “real” applications, however, is remarkably slow in some areas that could certainly benefit from it, such as in health care [6]. This is due to several reasons found both on the application and the technology “sides”, and, to a large extent, also due to a large gap between the communities that develop and use this technology. Filling or bridging this gap for the construction of usable and useful applications will determine whether grid technology will live up to its promises (or not).

In this paper we present and discuss attempts in the scope

of the Virtual Laboratory for e-Science (VL-e) project to bridge this gap for a particular application area, namely medical imaging for neuroscience research. The initial target application is functional Magnetic Resonance Imaging (fMRI), which enables the study of brain function. The goal is to develop a “Virtual Laboratory for fMRI,” or VL-fMRI, which will provide an IT infrastructure to facilitate management, analysis and sharing of data. The VL-fMRI offers resources for fMRI acquisition, data storage, computing and services, promoting the adoption of a meta-environment that serves as a virtual computer for the end-user. Access to local and remote resources is obtained with a simple, homogeneous, and friendly user interface of the Virtual Resource Browser (VBrowser).

The paper is organized as follows. In section 2 the structure of the VL-e project is introduced, and in section 3 the use case scenario for functional MRI is described. An overview of VLfMRI is presented in section 4, the VBrowser is presented in section 5, and the current status of the use case is presented in section 6. Related efforts are presented in section 7, and a discussion and final remarks close the paper (section 8).

2. Overview of the VL-e project Structure

The Virtual Laboratory for e-Science project (VL-e, www.vl-e.nl) brings together developers and users of grid technology with the goal of “boosting e-Science by creating an e-Science environment and carrying out research on methodologies.” The project consists of a Dutch consortium of academic and industrial partners formed in 2004. It is organized into subprograms that address different components of an e-Science environment: applications (e.g., medical, bioinformatics, biodiversity), virtual laboratory methodologies (e.g., workflow, parallel programming, databases, integration) and infrastructure (e.g., data and computing resources). The collaboration among part-

ners and subprograms is structured into “use cases” defined by the application areas. Working groups with representatives of all levels (application, virtual lab, and infrastructure) develop the use cases.

This paper refers to activities in the scope of one of the use cases defined by the “Medical Diagnosis and Imaging” subprogram. The goal of this use case is to provide an IT infrastructure to facilitate management, analysis and sharing of data in neuroimaging research, with a focus on functional MRI. This activity is carried out by a multi-disciplinary team of scientific programmers, system (integration) experts, (MRI) physicists and end-users affiliated with several departments of the University of Amsterdam (UvA), the Free University (VU), SARA Computing and Networking Services, and the National Institute for Nuclear Physics and High Energy Physics (NIKHEF). Expertise in complementary fields (IT, grid, networking, application development) is exchanged in regular meetings and permanent contact via a gforge project (www.gforge.org).

3. Use Case: Functional MRI

The cornerstone of this use case is the community of users of the scanning facilities provided by the AMC. Since 2003, the facilities include a 3 Tesla Philips MRI scanner (3T MRI, www.3tmri.nl), and in the near future a PET-CT scanner will also be available. The facilities are used for research and advanced clinical applications (e.g., neurosurgery planning).

Among other applications, the 3T MRI facility is shared by researchers that use fMRI to study brain function. Functional MRI is an imaging modality in which a subject is scanned while he/she is submitted to brain stimulation through physical or cognitive activity. The brain regions activated due to stimulation receive oxygen-rich blood, and the changes in the oxygenation level can be detected with MRI. The resulting scan consists of a series of 3D datasets (volumetric images) containing measurements along time, some of which obtained during stimulation and some at “rest”.

The scans, together with temporal stimulus data and other signals recorded during acquisition, are subsequently analysed in a series of steps that include temporal and spatial motion correction, normalisation, and statistical analysis. The analysis generates activation maps indicating the location of the regions where brain activation has been detected. Activation maps are used to characterize brain function in individuals (e.g., for neurosurgery planning) or in populations (e.g., for neuroscience research).

Figure 1 illustrates the acquisition and analysis workflow for one individual. From the raw scan to activation maps, the image analysis process can take around one hour (or more on simple desktop computers), and use/generate 1GB

of data per individual. For population studies, an additional statistical analysis step is performed based on the (aligned) activation maps of all individuals, generating a single activation map represented in a reference coordinate system, e.g., an average brain [7].




Figure 1. Data acquisition and analysis workflow for fMRI of an individual.

Currently, around 20 projects use the 3T MRI facilities for fMRI. The community of users consists of approximately 60 researchers distributed among several departments of the AMC (Radiology, Psychiatry), other institutes of the UvA (Psychology, Informatics), the Free University (VU Medical Center, Psychiatry, Psychology) and other Dutch institutes. This community gathers in bi-monthly meetings to discuss issues related to infrastructure, methodology and research.

Difficulties related to data management and analysis are recurrent topics in the fMRI user group meetings. There is a general wish for a more appropriate IT infrastructure to facilitate all types of tasks that must be performed in fMRI studies, as well as to facilitate sharing of data, resources and expertise. The goal of this use case is to take the first steps to build such IT infrastructure, which has been coined “Virtual Laboratory for fMRI,” or VL-fMRI. Although focused on fMRI at first, the infrastructure will offer general functionality and can be used in other scenarios, in a broader scope of studies that adopt medical imaging.

4. Overview of the Virtual Lab for fMRI

The main challenges for the construction of VL-fMRI have been discussed in detail elsewhere [13]. A subset of the challenges will be faced by pursuing the following goals:

- facilitate data gathering from dispersed and heterogeneous devices, by providing homogeneous access to the data recorded by the acquisition equipments;
- facilitate data storage and archival, by providing access to large capacity and long-term storage resources;

- enable high throughput for data analysis, by providing access to high performance computing resources to perform parallel analysis of mutually independent data;
- facilitate the data logistics and management in (large scale) fMRI studies via workflow automation;
- provide remote data access to the resources from workstations located anywhere;
- enable secure data sharing via access control mechanisms for users and groups; and
- facilitate data retrieval, by providing tools to generate and query metadata.

Figure 2 provides an overview of the VL-fMRI resources, which consist of acquisition devices, computing elements (CEs), storage elements, and services. The resources are provided by the VL-e Proof-of-Concept Environment (VL-e PoC) and managed by SARA and NIKHEF. Access to resources is granted to authorized users of the *vlemed* virtual organization (VO).

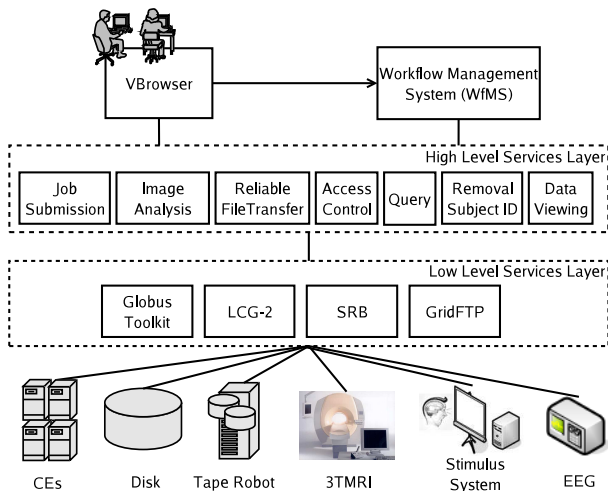


Figure 2. VL-fMRI resources accessed via the VBrowser: computing, storage and services.

The acquisition devices consist of a 3T MRI scanner, an EEG system, and a computer that controls the presentation of stimuli and instructions to the subject during the fMRI scanning session. These equipments are located at the scanning site (AMC), and connected to the intranet of the radiology department, which is protected by a firewall. Data (images, signals) can be exported from these devices via (secure socket layer) FTP. A special workstation, called

“grid access point” (GAP), is linked to the AMC intranet and the Internet, providing the necessary communication channel between these devices and the remote resources.

The CEs are accessed via the grid middleware of the EGEE project testbed (Enabling Grids for E-science in Europe, *public.eu-egee.org*). Currently, two clusters are available to the *vlemed* VO, with a total of around 140 worker nodes. Additional computing capacity will be available shortly with the installation of a new cluster at the AMC. This additional computing capacity will become available seamlessly via the EGEE middleware.

Storage resources consist of disks and tape silos for on-line and near-line storage space. These resources are managed with the Storage Resource Broker (SRB www.sdsc.edu/srb), which provides a seamless interface to store and retrieve federated data and metadata across a wide area network [3]. Control of access to the data and metadata can be performed at user or group levels, on a temporary or permanent basis. Currently the SRB server controls data that is stored in a centralized location (at the SARA), but in the near future other servers will be installed at the hospitals, in a true data federation setting. The file contents are transported automatically between the disk and tapes based on usage patterns (e.g., not accessed in the last N days). All files, on disk or tape, remain visible at all times for browsing, and the file content is automatically retrieved from tape when necessary, providing permanent and unlimited storage space. Data integrity and accessibility are guaranteed by automatic and periodic back-ups.

The high-level services layer implements functionality to perform varied tasks such as submission and monitoring of computing jobs, (fMRI) image analysis, reliable file transfer between local and remote file systems, data viewing, query based on metadata, access control, etc. These services are based as much as possible on existing applications, data formats and working environments that are already familiar to the end-users. For general tasks, services provided by the virtual laboratory methodology developed for the VL-e project are adopted. Services can be combined into scripts or composed into workflows with some Workflow Management System (WfMS), for example, the Data-Driven Workflow Engine (DDWE) [17]. The implementation of services follows a Service Oriented Architecture (SOA) approach, adhering to the Web Services Resource Framework (WSRF) standard.

All resources (acquisition, CEs, SRB, and services) are directly accessed by the end-users, who in this case are researchers with (typically) limited technical expertise. Currently available middleware focuses on the functionality rather than usability (e.g., command line interface), and as such does not qualify for direct usage by the target end-users of this application. This has been observed in the first phase of VL-fMRI, when only existing command-line tools were

used to upload/download data in the SRB and run jobs (see details in [13]). Not even the most advanced users could see a favorable cost-benefit ratio in that set-up.

There are several attempts of offering more friendly interfaces to grid resources, such as web portals (e.g. GridSphere [12]), interactive GUI-based clients (e.g., inQ SRB browser), and dedicated applications. These solutions aim at hiding the complexity underneath the implementation, sometimes at the cost of introducing simplifications that might constrain the available services or workflow configurations. Moreover, there is a clear separation between the local and the remote resources, usually implying large differences in the “look-and-feel” of the user interfaces that need to be operated by the users. In practice, the usage of grid resources remains a complex and difficult task for the end-users of this application domain, distracting them from the actual research topic to be addressed. This is perhaps also the case in other domains where users do not have sufficient technical background, or when the application requirements do not directly match those of many grid-related developments. To make the promises of e-Science reality, it is necessary to make grid resources available to non-technical users in more intuitive and transparent manners.

The VL-fMRI promotes the adoption of a meta-environment that serves as a virtual computer for the end-user, creating a single point of entry into “the Grid.” The idea is to provide access to local and remote resources with a simple, homogeneous, and friendly user interface, using a virtual desktop paradigm implemented by the VBrowser.

5. Virtual Resource Browser

The VBrowser is an interactive tool with a Graphical User Interface (GUI) that enables browsing local and remote (grid) resources from a single application. It offers an environment in which the user can interactively access resources of various types to manipulate data (upload, download, search, annotate, and view), run applications (prepare and run experiments) and monitor resources (status, control, notification). Its design is driven by requirements that emerged from several VL-e use cases in different application domains. Some of the requirements relevant in the VL-fMRI case are:

- Usability: easy to use, intuitive, and adaptable GUI, facilities for environment configuration;
- Single-point access to different types of local and remote resources, such as data storage, computation, and web services.

The adopted look-and-feel of VBrowser is inspired in GUI concepts from Microsoft Windows (Windows

XP), Linux (KDE/GNOME) and web browsers (Firefox/Netscape), creating a familiar and intuitive interface (Figure 3). The resource tree (left panel in Figure 3) shows the objects in each resource, for example, files and directories. Remote and local resources are handled in the same manner: the two directories in circle B in Figure 3 correspond to the local file system and to the SRB. Depending on the type of resource, the user is asked to authenticate with a password or with a GSI proxy (button in circle A in Figure 3) before browsing is enabled. Objects that belong to the resources are displayed as icons that indicate the object type. The user can manipulate the icons directly (right panel in Figure 3) to copy-paste, create links, delete, move, activate and view the content of objects. Multiple windows can be opened, and operations can take place in any (combination) of them. Configuration tools enable the user to install grid certificates, create profiles to access remote resources, and save (GUI) preferences.

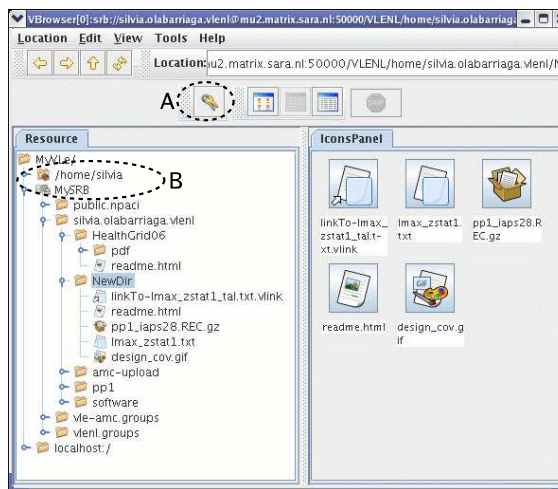


Figure 3. VBrowser GUI: resources panel (local and SRB file systems) and icons panel (files in the selected directory).

5.1. Resources

The VBrowser adopts a 3-tier architecture, implementing clients for several resources with a homogeneous GUI. A virtual resource layer provides abstract interfaces to resources of different types. The resources are part of the VL-e Toolkit, which contains a collection of generic services adopted and developed by the VL-e subprograms. The development new clients with a GUI to VL-e (web) services is facilitated by the virtual resource layer.

In the first development phase of the VBrowser, three resources are included: file management, reliable file transfer

and job execution.

Data Storage. A virtual data resource layer enables homogeneous access to storage devices in the local file system and in remote resources with the following protocols: GridFTP, Secure FTP and SRB – see Figure 4.




Figure 4. Virtual file resources.

Reliable File Transfer. The reliable file transfer (RFT) service provides a tool to move data among distributed grid-enabled storage resources. The RFT starts reliable third party file transfers and retries an arbitrary number of times in the case of failure. Permanent faults (if any) are notified to the user.

The RFT resource is based on the GT4 RFT service, which uses the GridFTP service and a relational database system to record the transfer-related events, enabling recovery in case of permanent failure. The RFT service can be activated from the VBrowser when the user schedules a file transfer, for example, with drag-and-drop between two (Grid) resources. The architecture is illustrated in Figure 5. The VBrowser communicates with the RFT Service, which manages a third party transfer between two storage servers (e.g., the SRB at the SARA and GAP node at the AMC). The Data Storage Interface (DSI), which is a GridFTP extension for SRB, is used to communicate with the SRB on behalf of GridFTP.

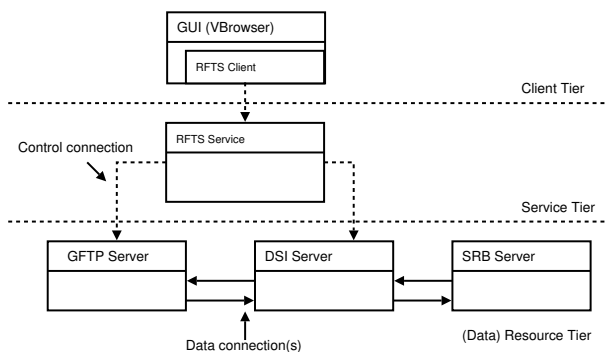


Figure 5. RFT resource.

Computing jobs. The job execution resource aims at providing access to computing resources from the virtual desktop to run medical imaging computing jobs. By manipulating icons associated with the computing resources (e.g., jobs and job schedulers), the user starts jobs, monitors and controls their status and execution, and inspects results. Special facilities enable managing the submission of very large number of jobs (job farming). Errors are informed to the user via a notification system (e.g., via GUI or e-mail). Interactive tools facilitate the design of new jobs.

At least two different types of job execution systems will be supported (see Figure 6): EGEE and DDWE [17]. In EGEE, all computation steps are captured into a script. Typical workflow steps are (1) download data and software from SRB into worker node; (2) run application; (3) upload results into SRB. The DDWE provides an interactive tool to build more complex workflows for processing data stored on geographically and grid-enabled distributed computation resources. Once the workflow has been composed interactively by the user, it is submitted to the workflow engine, which can then orchestrate its execution and monitoring. The data-driven workflow engine can submit multiple jobs on the available grid resources and create a direct connection between these jobs to allow data to flow among them. The workflow engine is implemented as WSRF-WS, and has the ability to select the appropriate computing resources based on information collected from both the application and the underlying grid infrastructure.

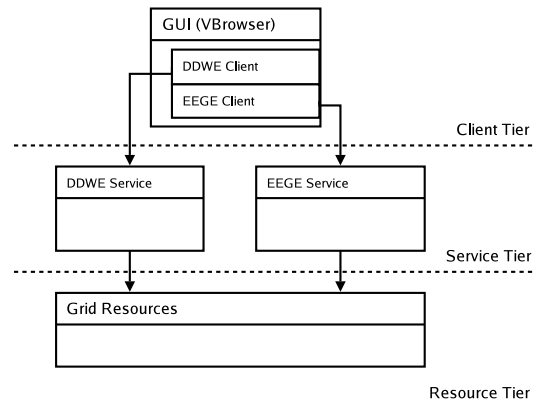


Figure 6. Job execution resource.

5.2. Implementation Aspects

The combination of the middleware functionality of the Globus Toolkit (version 4) and Java (1.5) provides a solid and broad platform for creating virtual environments. By using standard software engineering techniques (OOA/OOD, Object Oriented Analysis and Design) for the

creation of abstract (or virtual) classes, an abstract layer can be created on top of existing (grid) middleware. The virtual resource browser based on this virtual layer enables access to remote resources without the local operating system having to know about all of them, as long as the platform is available on the user's desktop. In time, standard drivers will be available at the OS level, and possibly the grid environment will become consolidated, such as the Internet and (local) area networking are today. For rapid prototype creation *today*, however, the Virtual Machine (VM) approach seems to be provide a proper solution.

The VBrower is implemented in Java, which provides a VM with an almost platform independent environment. The remaining dependencies are the version of Java and the Java Development Kit (jdk), which has a large set of libraries or classes that are directly useful to the developer. Additionally, the popular Swing framework offers almost directly usable GUIs. By using popular and readily available (web/grid) technology for Java, the construction of VBrower amounted to combining/assembling technologies such as grid file system or data storage APIs (SRB), MIME type association and viewers (e.g., HTML), web service integration, WSRF compliant Grid Service Integration, etc.

5.3. Current Status

The VBrower currently enables access to data resources (local file systems, GridFTP, Secure FTP and SRB) and reliable file transfer. A limited set of MIME types are supported: text, images (gif, png, jpeg) and html browsing (see Figure 7). The VBrower is adopted daily by around 5 beta-users of the VL-e project, in Windows and Linux desktops.

Plans for the near future include extending the range of supported MIME types (e.g., pdf, ps and viewing of medical images), enabling the activation of external applications for viewing (e.g., advanced visualization and spreadsheet software), and introduction of basic facilities for computing jobs.

6. Use Case Status

The basic infrastructure for data storage and job execution is already available and in use. The installation of the GAP at the AMC has been successfully completed after adequate configuration of firewall, network and OS settings. The VL-e PoC distribution software (*poc.vl-e.nl*) was valuable to set up the GAP system, facilitating installation and configuration of EGEE middleware and GT4 services. The GAP is now operational, and data can be transferred easily from the scanner to the SRB using the VBrower. A small group of advanced users currently perform a robustness test of the VBrower.

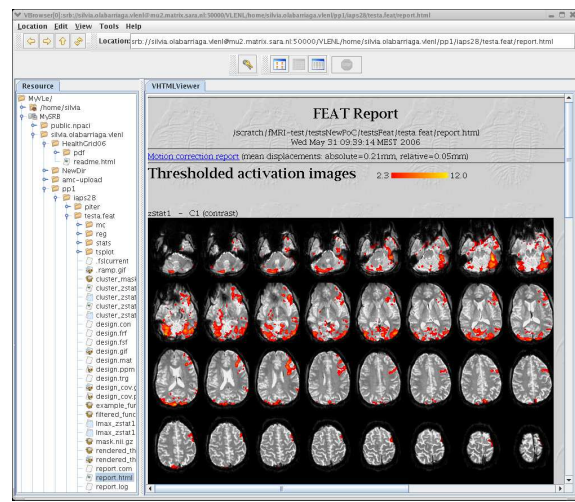


Figure 7. Browsing remote HTML files.

Computing jobs are still initiated and monitored via command-line utilities activated from scripts. A simple template workflow has been developed, facilitating the packaging of medical image analysis applications to run as jobs on grid resources. Current applications using job execution resources go beyond the scope of fMRI: non-rigid image registration for early detection of Alzheimer's disease; statistical discrimination of populations (schizophrenic patients and controls) based on Diffusion Tensor Imaging; and detection of activation sources in MagnetoEncefalography (MEG) data for epilepsy patients.

Although the (neuroimaging) applications above differ significantly in terms of software used and typical workflow, they all use the SRB as a file system to store the data manipulated in computing jobs. Intermediate results, progress information and final results of the computation are also uploaded into the SRB and can be directly inspected with the VBrower.

7. Related Work

The construction of grid-based systems for medical applications has been an active area of research and development - see overview of the possibilities and challenges in [6]. In some cases, the focus lies on the computation aspects: improvement of throughput or latency with parallel techniques, [10], construction and execution of workflows for image analysis (LONI Pipeline [14], MOTEUR workflow engine [9]), and image analysis services [15, 1]. In other cases, the focus lies on data aspects: integration of clinical information systems and grid middleware [11], access to distributed databases and information systems in the Mammogrid project [16], the NCR Informatics Initiative

[4] and the NeuroGrid project [8]. A few projects focus both on data and computing aspects, such as the Neurobase project [2], which uses grid technology for the integration and sharing of heterogeneous sources of information in neuroimaging.

In most examples mentioned above, the infrastructure has been validated, if at all, for a limited set of use case scenarios implemented by dedicated applications and portals. The reported developments are usually concentrated on functionality, and not on usability aspects. The VL-fMRI differs from the above initiatives because the goal is to address a large number of usage scenarios inserted in the workflow of a large number of researchers associated directly or indirectly to the 3T MRI scanning facility at the AMC. The development of custom software for each new use case would be prohibitive. Instead, a general and intuitive interface to access remote and local resources, the VBrowser, has been created for VL-fMRI. At the first sight, the VBrowser seems to be a regular browser, providing access to computing resources via a paradigm that is familiar to any user of desktop computers. Additional features, however, distinguish it from typical browsers. First, it is grid-enabled, using GSI authentication to access grid resources. Second, local and remote resources are manipulated in the same manner. Third, it is fully written in Java and (in principle) can run on a variety of platforms. And finally, its architecture allows for access to different types of resources from a single application and GUI.

8 Discussion and Concluding Remarks

This paper presented the initial steps taken in the VL-e project to build an IT infrastructure to facilitate management, analysis and sharing of data in neuroimaging research. The target application is initially fMRI, but the infrastructure is generic, serving not only other medical imaging problems, but potentially also other e-Science applications with similar requirements. The VL-fMRI implements a meta-environment approach that serves as a virtual computer for the end-user, creating a single point of entry to access local and remote resources. The VBrowser is the cornerstone of the user interface to the “Grid”, providing a simple, homogeneous, and friendly tool to reach and control all available grid resources.

According to the initial expectations, it should have been simple to build the VL-fMRI using existing tools, standards and protocols such as the Globus toolkit, SRB, EGEE middleware, portal builders, workflow management systems, etc. The pilot phase should have been completed in a few weeks, since it was just a matter of using existing grid technology develop by pioneering efforts in other application domains (e.g. High Energy Physics). Unfortunately, this assumption turned out to be incorrect due to several dif-

ficulties faced in the project, most of them also acknowledged by others before us [6, 12, 4]. Immature and fast evolving technology, shortage of trained professionals, an overwhelming amount and variety of overlapping and complementary software tools, unfriendly middleware, etc., are just some of the reasons why it turned out to be so difficult to implement VL-fMRI. The functionality is mostly available, but unfortunately the environment is still “hostile” and not directly usable by non-technical users.

The development of successful applications requires much effort, as well as the participation of experts with varied backgrounds. It is necessary to go beyond the “low hanging fruit” approach and invest as well on grid usability [5]. Tools such as the VBrowser seem to be a good start. Since its introduction (June/2006), it has received very good feedback from the users. The clean and intuitive user interface with facilities for direct browsing and content viewing have been decisive to attract reluctant users of grid technology in the medical subprogram of VL-e. After seeing a demo of the VBrowser functionality, one of such users immediately asked its installation in his desktop. We believe that the availability of additional functionality in the VBrowser (more advanced viewing tools and job submission) will be decisive for disseminating the culture of grid computing among the users of the VL-fMRI infrastructure.

Acknowledgements

This work was carried out in the context of the VL-e project, which is supported by a BSIK grant from the Dutch Ministry of Education, Culture and Science (OC&W) and is part of the ICT innovation program of the Ministry of Economic Affairs (EZ). We are grateful to many VL-e colleagues who shared with us their vision and expertise: Paul van Hooft, Jan Just Keijser, Dennis van Dok, Bart Heupers, Johan Alkemade, Robert Belleman. We also thank the directors of the AMC 3T MRI facilities, Ard den Heeten and Kees Grimbergen, and the director of VL-e, Bob Hertzberger, for his vision and support.

References

- [1] S. Bagnasco et al. Early Diagnosis of Alzheimer’s Disease Using a Grid Implementation of Statistical Parametric Mapping Analysis. In *Proceedings of HealthGrid*, pages 69–81. IOS Press, 2006.
- [2] C. Barillot et al. Federating Distributed and Heterogeneous Information Sources in Neuroimaging: The NeuroBase Project. In *Proceedings of HealthGrid*, pages 3–13, 2006.
- [3] C. Baru, R. Moore, A. Rajasekar, and M. Wan. The SDSC storage resource broker. In *CASCON’98 Conference*, 1998.
- [4] R. Begent et al. Challenges of ultra large scale integration of biomedical computing systems. In *Proceedings of the*

18th IEEE Symposium on Computer-Based Medical System, 2005.

- [5] F. Berman. Building a community grid: My view from the trenches. *GRID Today*, 1(6), 2002.
- [6] V. Breton, K. Dean, and T. Solomonides. The Healthgrid White Paper. In *Proceedings of HealthGrid*, pages 249–321. IOS Press, 2005.
- [7] A. C. Evans, D. L. Collins, , R. R. Mills, E. D. Brown, R. L. Kelly, and T. M. Peters. 3D statistical neuroanatomical models from 305 MRI volumes. In *Proceedings of the Nuclear Science Symposium and Medical Imaging Conference*, volume 3, pages 1813 – 1817. IEEE, 1993.
- [8] J. Geddes et al. NeuroGrid: Using Grid technology to advance Neuroscience. In *Proceedings of the 18th IEEE Symposium on Computer-Based Medical System*, 2005.
- [9] T. Glatard, J. Montagnat, and X. Pennec. Bridging clinical information systems and grid middleware: a Medical Data Manager . In *Proceedings of HealthGrid*, pages 93–103. IOS Press, 2006.
- [10] J. Montagnat et al. Medical Images Simulation, Storage, and Processing on the European DataGrid Testbed. *Journal of Grid Computing*, 2:387–400, 2004.
- [11] J. Montagnat et al. Bridging clinical information systems and grid middleware: a Medical Data Manager . In *Proceedings of HealthGrid*, pages 14–24. IOS Press, 2006.
- [12] J. Novotny, M. Russel, and O. Wehrens. GridSphere: A Portal Framework for Building Collaborations. In *Proceedings of 1st International Workshop on Middleware in Grid Computing*, 2003.
- [13] S. D. Olabarriaga, A. J. Nederveen, J. G. Snel, and R. G. Belleman. Towards a Virtual Laboratory for fMRI Data Management and Analysis. In *Proceedings of HealthGrid*, pages 43–54. IOS Press, 2006.
- [14] D. Rex, J. Ma, and A. Toga. The LONI pipeline processing environment. *NeuroImage*, 19:1033–1048, 2003.
- [15] V. Roberto et al. Neuroimaging Services on the Net. In *Proceedings of the 18th IEEE Symposium on Computer-Based Medical System*, 2005.
- [16] D. Rogulin et al. A grid information infrastructure for medical image analysis. In *Proceedings of DiDaMIC Workshop (MICCAI'2004)*, 2004.
- [17] D. Vasuinin, V. Korkhov, A. Belloum, and L. Hertzberger. Interactive dataflow driven workflow engine for grid enabled resources. *submitted*.