HDGPortal: A Grid Portal Application for Pervasive Access to Process-Based Healthcare Systems

Vassiliki Koufi and George Vassilacopoulos
Department of Digital Systems
University of Piraeus
Piraeus, Greece
{vassok, gvass}@unipi.gr

Abstract— **Healthcare is an increasingly collaborative enterprise, involving many individuals and organizations. The ability to provide readily access to integrated healthcare services remotely is now essential. Although Grid technology can support the integration of healthcare services across settings of care and among providers of care, it doesn't make these services available to authorized users at the point of care when needed. The integration of mobile and wireless devices with Grid technology can provide ubiquitous and pervasive access to Grid services. This paper presents HDGPortal, a Grid portal application which provides access to workflow-based healthcare processes using wireless Personal Digital Assistants (PDAs). HDGPortal's architecture is based on wrapping the workflow applications as web services and web service resources. These services are available to users through a number of Grid portal components that can be used for workflow enactment. HDGPortal has been implemented with a multi-layered security infrastructure in order to ensure secure access to healthcare processes.**

Keywords- Grid Portal, Web Services, Grid Database Services, Wireless Client, Pervasive Access, BPEL Engine, Personal Digital Assistant (PDA).

I. INTRODUCTION

Healthcare delivery involves a broad range of in-patient, out-patient and emergency healthcare services, performed by a number of geographically distributed and organizationally disparate healthcare providers [25]. As healthcare providers are mostly hosting diverse information systems, there is a need for an architecture that encompasses existing systems and aligns with the requirements of delivering shared and integrated care when and where needed. To this end, the architecture should focus on meeting the requirements of collaboration and coordination among healthcare services as well as the requirements of mobility of healthcare professionals (e.g. during ward rounds). Thus, patient information which is scattered around disparate and geographically dispersed systems can be readily accessed by authorized users at the point of care.

The incorporation of Grid technology in the healthcare domain led to the advent of Health Data Grids which serve as the integration infrastructure for shared and coordinated use of diverse data resources in virtual healthcare organizations comprised of geographically distributed and organizationally disparate healthcare providers [22][25]. Moreover, service orientation of the Grid makes it a promising platform for seamless and dynamic development, integration and deployment of service-oriented healthcare information systems using workflow technologies and Grid database services [10]. The latter can be created by using middleware products such as Open Grid Services Architecture - Data Access and Integration (OGSA – DAI) [12], which is part of the Globus Toolkit [28]. Grid database services are built as an extension to Web services and deliver added-value, high level data management functionality. For example, unlike traditional Web service solutions, OGSA-DAI supports dynamic creation of services which may combine and transform data from multiple distributed data resources (e.g. via a Distributed Query Processing service) in order to present an integrated or even derived view of the data [26]. As the existing Grid middleware, including Globus Toolkit, does not support mobility in the Grid nodes and is generally suitable only for high-profile service providers and clients [9], a possible extension of Health Data Grids may include the provision of access to the Grid database services in a ubiquitous and pervasive manner (i.e. via mobile and wireless clients).

Since their introduction in the early 1990s, Personal Digital Assistants (PDAs) have become increasingly popular for a large variety of applications in the healthcare domain [3]. A wireless network enabled PDA better meets a healthcare professional's needs because, unlike a portable PC, it slides smoothly in the gown pocket and is immediately available after startup, providing online access to different information services [11]. Due to the PDA's relatively low processor power and operating system restrictions, these services are mostly web-based. In turn, due to the low performance of the web page rendering engine in PDA browsers and the limited display size lightweight web-based solutions should be designed.

A Grid portal, being a secure and highly customizable online environment for gathering information about Grid services and resources, can deliver complex Grid solutions to users through a web browser without requiring to either download and install specialized software or set up networks, firewalls and port policies accordingly [8][19]. Hence, Grid-enabled portals have been proven to provide effective mechanisms for exposing distributed systems to general user communities without forcing them to deal with the

PERVASIVEHEALTH 2008, 30 Jan - 1 Feb. Tampere, Finland

DOI 10.4108/ICST.PERVASIVEHEALTH2008.2519

complexities of the underlying systems [19]. On these grounds, a Grid portal application that complies with the restrictions imposed by PDA technology seems an ideal solution for the provision of pervasive access to Grid healthcare processes.

Along these lines, we present HDGPortal, a Grid portal application that provides a web interface to workflow-based healthcare processes spanning a health district. HDGPortal provides pervasive and ubiquitous access to healthcare processes which are modeled as flows of Grid database services that expose the data resources, such as relational and XML databases, residing within a health district. The aforementioned processes may involve access to medical data by authorized users or triggering a new healthcare process, such as ordering of laboratory tests.

HDGPortal is designed for being accessed by PDAs. Security in portal transactions is provided by the Grid Security Infrastructure (GSI), which is based on Public Key Infrastructure (PKI). Additional security mechanisms, regarding access control are also utilized in order to ensure authorized access to healthcare processes through HDGPortal Grid portal application.

II. Related Work

During the last few years, research efforts regarding the advancement of pervasiveness of Grid Computing have been on the rise. Some involve the provision of pervasive and ubiquitous access to existing Grid infrastructures, by enhancing the scope of present Grid environments to the emerging mobile domain. The GridLab project [7] proposes an architecture for mobile clients whereby J2ME-enabled mobile devices are incorporated as the clients of Grid services to enable Grid users access their applications and resources remotely [6]. However, the proposed architecture focuses on scientific computations and increases the overhead on mobile devices.

A pervasive Grid architecture which allows low profile clients to access Grid database services in a pervasive manner is proposed in [9]. The proposed architecture is intended for use through mobile devices and its main focus is to support day-to-day data-intensive query applications that are used in the scientific world as well as in the commercial world. Although this architecture provides seamless access to Grid database services, it doesn't provide access to workflows of such remote services. A similar approach, with the same deficiency, is presented in [15], whereby integrated view of services related to data management is provided in a pervasive Grid. A drawback of this approach is the limited ease of distributing and extending the software platform developed for hosting pervasive services (PerSE [18]), since it does not rely on well established standards.

The Genius Grid Portal [5] is a problem solving environment based on a middleware developed by the European EGEE Project [30] and provides a high level graphic interface for many applications belonging to different scientific areas, including bio-medicine. However, lack of support for several data management services (it only supports replication services) constitutes it unsuitable for use in the healthcare domain.

In addition to the aforementioned studies, a considerable amount of research effort has been devoted to tackling various aspects of pervasive systems in the area of Grid computing [1][4][14][20][21], but none of them handles data management and access control in the context of a process-oriented healthcare environment.

In this paper, we propose a Grid portal application that provides pervasive and authorized access to healthcare processes. These are modeled as Business Process Execution Language (BPEL) processes that orchestrate geographically dispersed Grid database services concerned with healthcare delivery within and across healthcare organizations [24]. The Grid database services are created using OGSA-DAI middleware to provide integrated access to healthcare databases irrespective of the network location they reside.

III. Motivating Scenario

To illustrate the main principles of the proposed system architecture and implementation, a sample integration project is described which is concerned with the automation and remote access to cross-organizational healthcare processes spanning a health district. Typically, a health district consists of one district general hospital (DGH) and a number of peripheral hospitals and health centers. As patient referrals are usually made among various healthcare providers within a district (e.g. for hospitalization, for outpatient consultation or for performing specialized medical procedures), there is a need to ensure that automated healthcare processes can be executed remotely and integrated patient information can be accessed through the execution of these processes by authorized users where and when needed. Thus, the sample process considered here is concerned with patient referrals issued by healthcare professionals, while on a ward round, to a clinical department of the same or another hospital. This process involves two different operations underpinned by Grid and workflow technologies.

- Healthcare process execution: an operation that involves handling medical order processing among health district organizations that is exposed as a workflow of Grid database services.
- Medical record data access: an operation that involves patient medical data retrieval from dispersed healthcare organizations that is exposed as a Grid database service (i.e. an OGSA-DAI distributed query, which enables execution of queries in parallel over OGSA-DAI database services and other Web services on the Grid).

Suppose a healthcare delivery situation where a hospital's physician, while on a ward round, visits one of his/her patients. While assessing the patient's condition at the point of care, the physician may need to access the patient's medical record and to issue a request for a set of laboratory tests. After receiving the order, the laboratory schedules the tests requested and sends a message to the requesting physician notifying him/her on the date and time scheduled for the specimen collection and test performance.

IV. PORTAL ARCHITECTURE

Figure 1 shows a high-level system architecture, which is described by a three-tier model, comprising of the PDA client, the server site of the DGH and the Grid.

The first tier is the PDA client. The PDA is running an HTTP(S)-based client, which is the PDA's web browser and provides user interaction with the system.

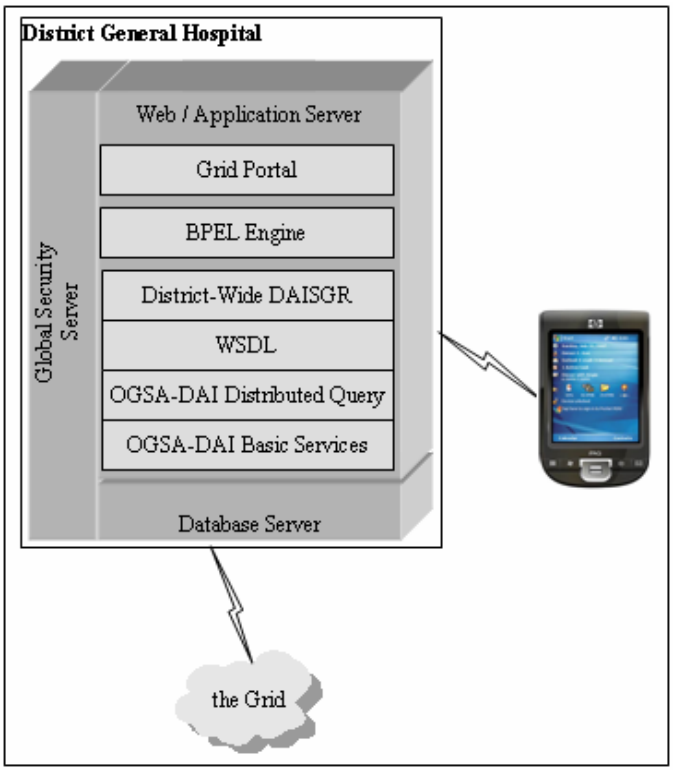

Figure 1. System Architecture

The second tier, which essentially constitutes the HDGPortal's environment, is hosted on a server at the DGH site and consists of the following components:

- *OGSA-DAI middleware*: it is the middleware product used for the generation of the specific Grid database services that manage access to district-wide medical data. These services may be created by using either OGSA-DAI Basic Services, for access to DGH's database, or OGSA-DAI Distributed Query for retrieval of district-wide consolidated patient information. A district-wide registry, DAISGR, is used for discovery of the database services required for a task. WSDL provides the description of these services and how they can be accessed.
- *BPEL Engine*: it handles the execution of BPEL-based healthcare processes provided in the healthcare settings within a health district. The role of the BPEL engine is, given a BPEL process definition and a set of inputs, to instantiate it, executing the tasks calling the various Grid database services and routing the data between them.
- *Grid Portal*: it provides a web-based front end to Grid resources. It consists of a JSR-168 [32] compliant portlet container that hosts and manages the HDGPortal BPEL portlet as well as the portlets of the workflow (BPEL) applications. A portlet is a java web component that generates dynamic content in response to processed requests [33]. The HDGPortal BPEL portlet provides the Web browser-based portal user interface to the BPEL Engine where all healthcare processes are deployed. Thus interaction with the corresponding workflows is enabled. This interaction is performed through certain portlets developed to facilitate the physicians' interaction with the relative tasks of the workflows.
- *Web/Application Server:* It provides the hosting environment to the aforementioned components.

The third tier of the system architecture is the Grid (i.e. the Health Data Grid) which comprises of remote data resources. These are heterogeneous and reside in geographically distributed and organizationally disparate healthcare providers within a health district. Each one of the Grid nodes residing at a healthcare setting has the structure illustrated in Figure 2.

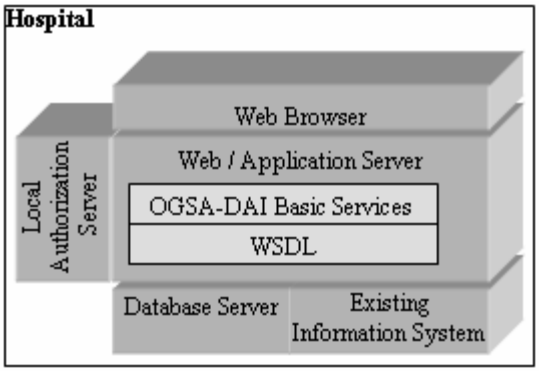

Figure 2. Hospital Grid Node

The database used for the storage of medical information in each healthcare setting is Grid-enabled and the database services exposing it to the Grid are deployed in the web server of the organization and registered with the district-wide DAISGR.

A. HDGPortal BPEL Portlet

The BPEL portlet is a java-based client application to the BPEL Engine. The portlet uses the engine's web service in order to provide the required functionality to the users. In particular, a physician accesses the BPEL portlet to instantiate a BPEL process and execute the tasks he/she is authorized to access. On task execution, the associated Grid database services are located and invoked to update patient records or retrieve district-wide consolidated patient information.

B. Security Issues

One critical issue in a pervasive healthcare system build on top of a Grid infrastructure is security. To this end, HDGPortal implements a multi-layered infrastructure for authentication and access control services.

With regard to authentication, HDGPortal utilizes both a Lightweight Directory Access Protocol (LDAP) directory service and Grid Security Infrastructure (GSI) authentication mechanism, which is part of the Globus project [28].

LDAP directory is used for user/password based authentication of users (i.e. healthcare professionals) to HDGPortal. LDAP is a directory service protocol based on the client-server model and runs over TCP/IP [31]. The LDAP server is hosted at the DGH and contains information about all the users authorized to access HDGPortal. This information include the healthcare professionals' username and password.

GSI authentication mechanism is used for Grid authentication and defines single sign-on algorithms and protocols, cross-domain authentication protocols, and temporary credentials called proxy credentials stored in the MyProxy Server [13]. MyProxy is open source software for managing X.509 Public Key Infrastructure (PKI) security credentials (certificates and private keys) [13]. It combines an online credential repository with an online certificate authority to allow users to securely obtain proxy credentials when and where needed, without worrying about managing private key and certificate files [13]. In our environment, a MyProxy server is hosted on a server at the DGH site. Healthcare professionals use MyProxy to delegate credentials to the HDGPortal, which then is acting on their behalf. This is achieved by storing credentials in the MyProxy repository and sending the MyProxy passphrase to the HDGPortal.

All web transactions are executed under the Secure Socket Layer (SSL) via HTTPS.

Authorization and access control constitute major research issues both in the traditional and pervasive Grid environments. In our environment, access control is provided at two levels: the BPEL task level and the Grid database service level. Access control is one of the aspects of business processes explicitly mentioned to be outside the scope of BPEL [23]. Thus, a BPEL process can be executed by everybody who is able to support the relevant partner link types [23]. However, this is inappropriate in the healthcare domain due to the strict security requirements on medical data managed by the database services incorporated into the BPEL process tasks. Therefore, when a user logs in the HDGPortal to initiate an interaction with BPEL Engine in order execute a BPEL process, a role-based access control mechanism is used to determine whether access to certain BPEL tasks should be permitted or denied. Thus, RBAC is enforced into each BPEL process in order to meet the access control requirements of the healthcare domain under study.

Access control at the Grid database service level is enforced through the use of a Community Authorization Service (CAS) Server. This is hosted on a server at the DGH site and provides a scalable mechanism for specifying and enforcing complex and dynamic policies that govern data resource usage within the Health Data Grid [22]. The use of CAS, in conjunction with the local Access Control Lists (ACLs) held at each healthcare setting, establishes Role-Based Access Control (RBAC) within the OGSA-DAI framework [22].

V. PROTOTYPE IMPLEMENTATION

To illustrate the functionality of the proposed architecture, we describe a prototype system which is based on the case scenario of Section III. A high-level representation of the system functionality is depicted in Figure 3.

The prototype implementation of the proposed Grid portal application and the underlying security services has been developed on a local Grid.

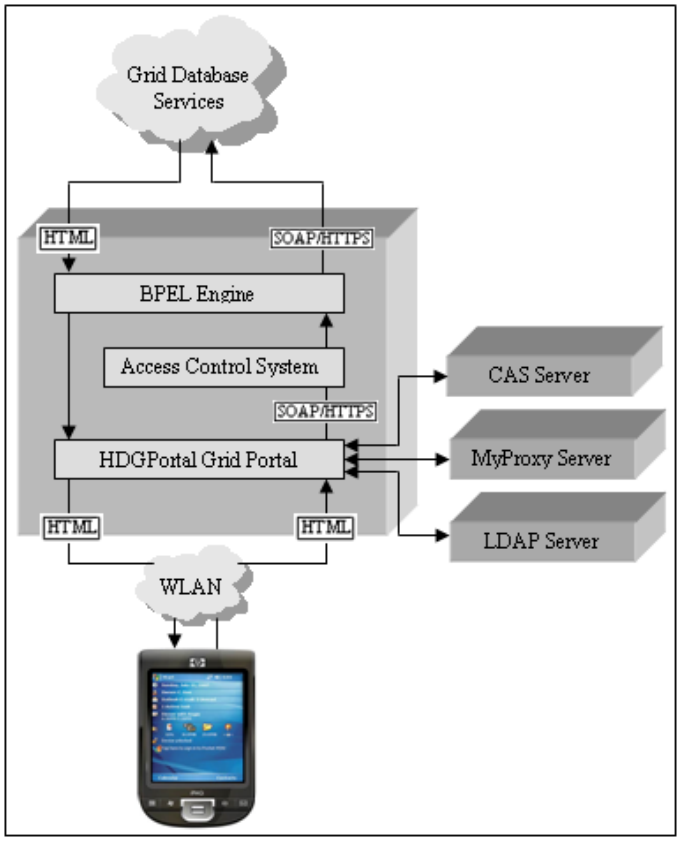

Figure 3. System Functionality

The software tools used are Apache/Tomcat as Web/Application Server and Globus Toolkit (GT4) for the provision of Grid middleware services. OGSA-DAI and its extension, OGSA-DQP, are used for the generation of the Grid database services using the geographically dispersed data resources. In our implementation, a PostGres SQL relational database and a MySQL database are used as data resources. The port types and descriptions of the Grid database services are described by generated WSDL, which also indicates the way these can be accessed. The BPEL engine used for the execution of BPEL healthcare processes is ActiveBPEL [29], an open source BPEL Engine [2]. HDGPortal that provides access to this engine is based upon GridSphere portal framework [34]. The GridSphere Portal is a Java Servlets/JSP based framework that builds upon the Java CoG Kit [27], Globus Toolkit [28] and MyProxy [13] to support such features as single sign-on, job submission and data movement. The GridSphere portal framework is based upon the Java Portlet model [6]. The portlet concept is a Java standard (JSR-168) [32].

At the beginning of a shift, a physician, using his/her X.509 personal certificate and its private key, creates a temporary certificate (the Proxy Certificate), which is stored on the MyProxy Server and is valid for the duration of his/her shift. Then, the physician receives access to the Grid portal and, consequently, to all the healthcare processes he is authorized to access.

While on a ward round, the physician logs in HDGPortal using his/her LDAP username and password. HDGPortal contacts the MyProxy Server and requests the physician's Proxy Certificate from the MyProxy online credential repository. Hence, the physician's temporary credentials are retrieved, and passed on to the CAS server, which returns a new credential, known as a CAS proxy credential that contains the CAS policy assertions representing the physician's capabilities (in the form of roles) and restrictions as an extension [22]. Before an interaction with the BPEL engine is initiated, a role-based access control mechanism is invoked, whereby permission or denial of access to certain BPEL tasks is determined in accordance to the policy assertions contained in the physician's CAS proxy credential.

Suppose that a physician is permitted to execute a task concerning the retrieval of the patient's medical record. This task may involve data retrieval using a number of remote Grid database services that access the relevant data resources of the various healthcare settings where the patient has received care. For each remote data resource involved, the portlet authenticates itself (achieving single sign-on through its delegated proxy certificate) and receives the authorized privileges in order to execute the relevant Grid database service on behalf of the physician. To this end, the roles extracted from the extension of the CAS proxy certificate are checked against the local access control list. If the physician has the required privileges, the service is executed and the output, which is in XML format, is transformed in HTML format and rendered back into the physician's web browser. This transformation is supported within the OGSA-DAI framework through XSL Transform. For each transformation a corresponding style sheet is selected from a predefined set of XSLT style sheet documents. There is one style sheet for each document type (e.g. discharge letter, structured laboratory result, narrative text like radiology results, etc.) for the device type the physician uses (i.e. the PDA).

After reviewing the patient's medical record, a physician may order a set of laboratory tests by initiating the corresponding BPEL process and executing the relevant task.

VI. Concluding Remarks And Future Work

Grid workflow systems can offer great benefits in the development of healthcare information systems. However, for these systems to reach their full potential in providing quality and effective medical services, access to healthcare processes as well as to updated patient medical records anytime, anywhere must be provided to healthcare professionals. Although Grid technology supports the integration of healthcare services, currently it does not support the access of Grid database services in a pervasive and ubiquitous manner. A solution to this problem can be provided by incorporating mobile and wireless devices as clients of Grid services. However, due to certain limitations of these devices (e.g. as small display size), a need for a lightweight data Grid client arises in order to facilitate easy and quick access to medical data and services. The HDGPortal presented in this paper is a portlet-based application that provides a web interface to workflow-based processes spanning a virtual healthcare enterprise. This interface provides pervasive and ubiquitous authorized access to Grid healthcare processes and it is as friendly as the ones developed for use through desktop PCs whilst conforming to the limitations and restrictions of a PDA environment.

Security issues, such as access control and security of data transfer over the wireless network, suggest directions for future work. With regard to access control, a context-aware access control service could be considered as an alternative mechanism to enable support for granting and adaptation of permissions to healthcare professionals according to the current context. For example, the access privileges of a physician using HDGPortal on his/her PDA, should adjust depending on his identity, location and time of access as well as network status. As far as security of data transfer is concerned, a possible candidate among existing solutions for our environment could be the development of a Virtual Private Network (VPN).

References

[1] S. R. Egglestone, M. N. Alpdemir, C. Greenhalgh, A. Mukherjee, and I. Roberts, "A portal interface to myGrid workflow technology", All Hands Meeting (AHM2005), 19-22 September 2005, Nottingham.

[2] W. Emmerich, B. Butchart, L. Chen, B. Wassermann, and S. L. Price, "Grid service orchestration using the Business Process Execution Language (BPEL)". J. Grid Comp 2006; 3: 283-304.

[3] S. Fischer, T. E. Stewart, S. Mehta, R. Wax, and S.E. Lapinsky, "Handheld computing in medicine", J Am. Med. Inform. Assn 2003; 10(2): 139-149.

[4] D. Gannon, J. Alameda, O. Chipara, V. Dukle, L. Fang and M. Farrellee, "Building grid portal applications from a web-service component architecture", Invited Paper, Proceedings of the IEEE, vol. 93, no. 3, March 2005.

[5] Genius Portal, https://genius.ct.infn.it/

[6] P. Grabowski, B. Lewandowski, and M. Russell, "Access from J2ME-enabled mobile devices to Grid services", White Paper.

[7] GridLab, "A Grid application toolkit and testbed," http://www.gridlab.org

[8] F. Juan, J. Fox, and M. Pierce, "Grid portal system based on GPIR", Proceedings of SKG2006 Guilin China, October 31 - November 3, 2006.

[9] S. Khatua, S. Dasgupta, and N. Mukherjee, "Pervasive access to the Data Grid". Proceedings of the International Conference on Grid Computing and Applications, 26-29 June 2006, Las Vegas, Nevada, USA.

[10] V. Koufi, D. Papakonstantinou, and G. Vassilacopoulos, "Virtual patient record security on a Grid infrastructure". International Conference on Information Communication Technologies in Health (ICICTH'06), 13-15 July 2006, Samos, Greece.

[11] M. Muller, T. Frankewitsch, T. Ganslandt, T. Burkle, and H. U. Prokosch, "The Clinical Document Architecture (CDA) enables Electronic Medical Records to wireless mobile computing", MEDINFO, pp 1448-1452.

[12] The OGSA-DAI project, http://www.ogsa-dai.ac.uk/

[13] MyProxy Credential Management Service, http://grid.ncsa.uiuc.edu/myproxy/

[14] M. Parkin, and J. M. Brooke, "A PDA client for the computational Grid". Concurr. Comp. - Pract. E, 2007;19:1317 -1331.

[15] J. M. Pierson, “A pervasive Grid, from the data side”. Research Report, LIRIS Laboratory, 2006.

[16] K. Raatikainen, H. B. Christensen, and T. Nakajima, “Application requirements for middleware for mobile and pervasive systems”. Mobile Computing and Communications Review, 2002; 6(4):16-24.

[17] A. Ranganathan, and S. McFaddin, “Using workflows to coordinate web services in pervasive computing environments”. Proceedings of the IEEE International Conference on Web Services (ICWS’04), 6-9 June 2004, San Diego, California, USA.

[18] V. M. Scuturici, J. M. Pierson, and L. Brunie, “Perse : Pervasive services environment”. Research Report, 2005.

[19] M..P. Thomas, J. Burruss, L. Cinquini, G. Fox, D. Gannon and L. Gilbert, “Grid portal architectures for scientific applications”. J. Phys. Conf. Ser. 16: 596-600.

[20] B. Wu, M. Dovey, M. H. Ng, K. Tai, S. Murdock, P. Jeffreys, S. Cox, J. Essex, and M. S. P. Sansom, “A Web / Grid portal implementation of BioSimGrid: A biomolecular simulation database”. Proceedings of the International Conference on Information Technology: Coding and Computing (ITCC’04).

[21] I. Maglogiannis, A. Chatzioannou, J. Soldatos, V. Milonakis and Y. Kanaris, “An application platform enabling high performance Grid processing of microarray experiments”, 20th IEEE International Symposium on Computer-Based Medical Systems (CBMS 2007), 20-22 June 2007, Maribor, Slovenia.

[22] A. L. Pereira, V. Muppavarapu, and S. M. Chung, “Role-Based Access Control for Grid database services using the community authorization service”. IEEE T. Depend. Secure, 2006;3(2):156-166.

[23] J. Mendling, M. Strembeck, G. Strembeck, and G. Neumann, “An approach to extract RBAC Models from BPEL4WS processes”, Proceedings of the 13th IEEE International Workshops on Enabling Technologies: Infrastructure for Collaborative Enterprises (WETICE'04), 14-16 June 2004, Italy.

[24] A. Ranganathan and S. McFaddin, “Using workflows to coordinate web services in pervasive computing environments”. Proceedings of the IEEE International Conference on Web Services (ICWS’04), 6-9 June 2004, San Diego, California, USA.

[25] F. Malamateniou and G. Vassilacopoulos, “Developing a virtual patient record using XML and web-based workflow technologies”. Int. J. Med. Inform., 2003; 70(2-3): 131-139.

[26] M. Antonioletti, N. C. Hong, A. Hume, M. Jackson, A. Krause and J. Nowell, “Experiences designing and implementing Grid database services in the OGSA-DAI project”, Designing and Building Grid Services Workshop, GGF9, October 2003.

[27] Java CoG Kit, http://www-unix.globus.org/cog/

[28] The Globus Toolkit, http://www.globus.org/

[29] Active Endpoints, ActiveBPEL Open Source Engine Project, http://www.activebpel.org/

[30] EGEE (Enabling Grids for E-sciencE) Project, http://www.eu-egee.org/

[31] S. AbdelMalek, LDAP an Overview, http://www.eecg.toronto.edu/~jacobsen/courses/mwci/scribes/ldap/LDAP.htm

[32] JSR-168 Portlet Specification, http://www.jcp.org/aboutJava/communityprocess/final/jsr168/

[33] D. Del Vecchio, V. Hazlewood and M. Humphrey, “Evaluating Grid portal security”. Proceedings of Supercomputing, November 2006, Tampa, FL.

[34] Gridsphere Portal Framework, http://www.gridsphere.org/gridsphere/gridsphere