MANAGING VIRTUAL ORGANIZATION PROCESSES BY SEMANTIC WEB ONTOLOGIES

A. Gehre, P. Katranuschkov, R. J. Scherer
Institute of Construction Informatics, Technical University of Dresden, Germany

ABSTRACT: Interoperability within Virtual Organisations (VOs) is still only weakly supported by IT frameworks. Whilst service level interoperability has made remarkable progress since the emergence and the rapid growth of SOA and Grid technology in the last years, business processes – which are the driving force of each VO – still suffer distinct conceptual gaps regarding their decomposition to technical transactions. There exists no detailed approach that would allow describing technical as well as business aspects in a coherent yet flexible and extensible way. This paper presents a newly developed semantic framework that targets this requirement. The conceptual background is followed by an introduction of the developed semantic web ontologies. Based on these definitions, dedicated Ontology Services as well as a set of related end-user applications facilitating semantic technology have been designed and implemented. They are presented in the second part of the paper. Reported are results from the EU project InteliGrid (IST-004664; 2004-2007).

KEYWORDS: semantic web, ontologies, virtual organisation, process modelling, process management.

1 INTRODUCTION

The term *Virtual Organisation* (VO) is used widely today to denote inter-company business collaboration focused on a specific project context. However, the essence of this collaboration is defined variably by different research domains concerned with the VO concept. Within the Grid domain, a VO is considered as an entity formed by a dynamic multi-institutional set of individuals and/or institutions establishing common rules for coordinated resource sharing and problem solving. The focus lies not on file-sharing, but rather on “direct access to computers, software, data, and other resources, as is required by a range of collaborative problem-solving and resource-brokering strategies emerging in industry, science and engineering. This sharing is, necessarily, highly controlled, with resource providers and consumers defining clearly and carefully just what is shared, who is allowed to share, and the conditions under which sharing occurs” (Foster et al. 2001). From the business perspective, a VO is defined differently. According to (Camarinha-Matos & Afsarmanesh 1999), a VO is “a temporary alliance of enterprises that come together to share skills or core competences and resources in order to better respond to business opportunism, and whose co-operation is supported by computer networks”.

Thus, the grid domain definition of a VO is about computer networks, resources and access rights, whereas the business-centred definition of VOs considers these issues as necessary technical support, while focussing on business aspects of VOs like business processes and opportunities. However, an integrated grid-based VO environment can only be successfully deployed and established on the market, if both aspects are coherently merged, i.e. the technical power of modern grid environments and instant support for integrated business processes.

What is missing for achievement of this level of integration is a *Semantic Framework* that can mediate between grid technology and the real business layer of the VO, thereby providing sustainable semantic support for VO management and integrated collaboration processes. This semantic layer has to deal with a broad spectrum of conceptual and technical requirements, as defined and ranked in (Turk et al. 2006). Subsumed from these requirements, a semantic framework for Grid-based VO environments has to:

- provide support for VO organisational structuring, including dynamic reorganisation during VO lifetime, as partners can leave and join at any time;
- support the organisational polymorphism that is typical for VOs;
- provide generic description, addressing and accessing of VO resources (data and services), including directory services with advanced semantic search mechanisms;
- enable business process integration on semantic level;
- support different levels of granularity, i.e. it should allow for hierarchical composition and decomposition of concepts;
- be extensible in order to react to changes, extensions and new business opportunities;
- ensure privacy of information;
- provide services that keep the information about the collaborative VO network up to date;
- ensure reasonable response times and scalability to ensure end user acceptance;
- use or provide for compatibility with industry standards.

The benefits of a semantic framework and the purpose of ontologies for its achievement are explained in a number of recent publications (cf. Berners-Lee et al. 2001, Miller 2003, Herman 2004). The hierarchical position of ontology specifications and their inter-relationship to other specifications related to the World Wide Web is best illustrated by the so-called Semantic Web Stack (Herman 2004).

An important prerequisite for achieving semantic interoperability is the ability to make use of rich computer languages, such as the Web Ontology Language OWL (W3C 2004a) that allows to describe the various entities in the VO environment in a coherent way. OWL has been developed recently on top of the existing XML, RDF and RDFS standards (W3C 2004b) which were not found sufficient for achieving adequate semantic interoperability. Although XML DTDs and XML Schemas seemed satisfactory for exchanging data between parties who have previously agreed to some set of shared definitions, their lack of constructs to describe the deeper meaning of these definitions prevents computers from reliably performing such tasks. This is explained in more detail in (Gehre et al. 2005), including a short discussion of available tools and frameworks for management of semantic data, as well as a list of projects also trying to apply semantic technology and ontologies for collaboration purposes.

2 THE ONTOLOGY FRAMEWORK CONCEPT IN A NUTSHELL

On the basis of the above considerations, an ontology framework for semantic VO interoperability has been designed and implemented in the EU project InteliGrid (IST-004664). Figure 1 below shows the overall concept of the framework together with its intended use in the VO environment. On the right hand side, it presents schematically the envisaged distributed run-time environment using the semantic grid paradigm. Users of the environment connect to it via their ontology-based virtual desktops that provide access to all environment services. On the left hand side, the figure presents the principal components of the ontology framework providing the basis for achieving semantic interoperability and collaboration in the environment.

It is important to understand that concepts defined in the ontologies on the left side of the figure act as *semantic proxies* of entities in the "real" environment. They allow capturing semantic metadata about the "things" in the distributed environment in a coherent way, although separated from the real entity itself. Since the full ontological expressiveness can be used there, complex relational information can also be captured, making coherencies *explicit* that are only implicitly contained in the "real" environment. For example, advanced classifications can be provided for information and service resources, files that are available on different locations with different access methods can be pooled in a singe file resource instance with references to the different locations and respective access methods, various cross-model relationships can be defined and processed, actor roles can be subsumed on the basis of only a few given attributes etc.

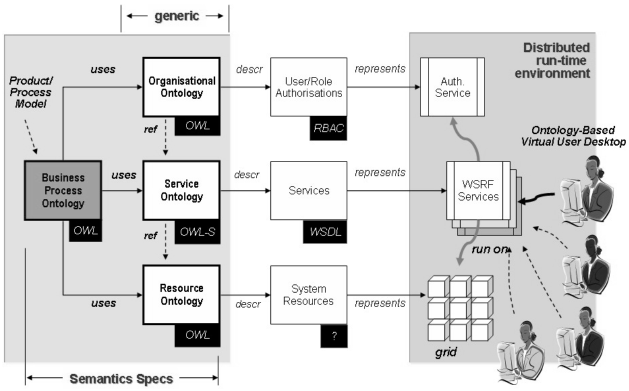

Figure 1. Overview of the developed ontology framework for semantic VO interoperability.

Due to various reasons (already existing specifications and services, security and performance considerations etc.) run-time components cannot be expected to be fully conformant with the ontology specifications comprising the semantic framework. Therefore, in the middle part of the figure typical mapping components are shown that have the task to "translate" ontology concepts to existing practical schemas and data structures as e.g. the RBAC model (role based access control, see Ferraiolo et al. 2001) defining access rights and authorisation of users. Whilst at this level many different variations are theoretically possible, current technologies appear to converge to a relatively small set of specifications. This makes the mapping task manageable in the context of specific industry contexts.

Figure 2 below reproduces the general schema of Figure 1 with focus on the ontology specifications and some of the defined top-level concepts. The shown relationships between the individual ontology categories indicate the principal dependencies and intended usage.

The *Resource*, *Services* and *Organisational* ontologies only reference each other as shown, but are not dependent on the *Business Process Ontology*. Therefore, they can also be used independently, providing basic semantic descriptions related to organisational entities, resources and services that can be implemented within a generic *Ontology Service*. In contrast, the Business Process Ontology is massively dependent on these lower level core ontologies. This allows it to provide more complex semantic support.

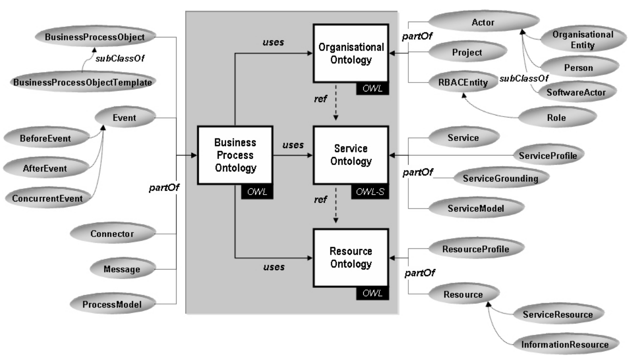

Figure 2: Selected ontology concepts

The developed modular modelling approach provides for different levels of support and different configurations of the environment for each specific case. It also greatly facilitates the modular development of the respective ontology services by clearly defining their information and functionality boundaries.

All ontologies are modelled in OWL-DL depth using the Protégé Ontology Editor. As an example, Figure 2 below shows the top level class concept “Resource”. Datatype properties are depicted in white ovals, class concepts in dark ones. As can be seen, services do also have a definition in the *Resource Ontology*. The “*ServiceResource*” concept, a direct child concept of “*Resource*” provides metadata definitions for describing services conformant to the Resource Ontology and a reference to in-depth service descriptions based on the OWL-S *Service Ontology*.

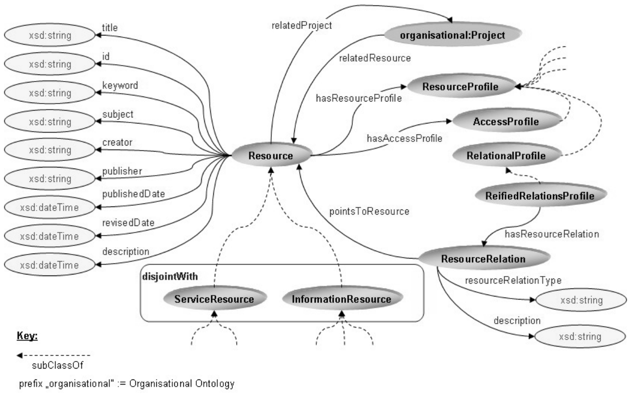

Figure 2. Selected classes and properties assigned to the core Resource concept.

More details about the methodology and the concepts of the modelled ontologies are provided in (Gehre et al. 2006). Following the semantic web approach, the ontologies are deployed on a public web server at *http://cib.bau.tu-dresden.de/ontologies/InteliGrid/generic/,* and can thus be easily referenced there by any future extensions and cross-linked ontologies.

OWL-S is used as is. More information about OWL-S can be found for example in (Martin et al. 2005).

3 ONTOLOGY SERVICES

Developing ontologies for semantic VO interoperability is only the initial step in supporting real business cases as targeted by the InteliGrid platform. For the success of an ontology-based approach, it is of utmost importance to implement *ontology services* that provide convenient methods for management of ontology instances, i.e. semantic metadata about entities in the IT environment. As mentioned before, these instances act as semantic proxies with descriptive information about entities like resources and users represented in the environment.

In the InteliGrid architecture, the *Ontology Services*, including a service for management of *Business Process Objects* (BPO Service), constitute the layer "Semantic Interoperability Services", as shown in Figure 3 below. They are connected to the "Business Services" layer and make use of the "Grid Middleware Services" that provide basic authorisation management and generic access to all grid resources.

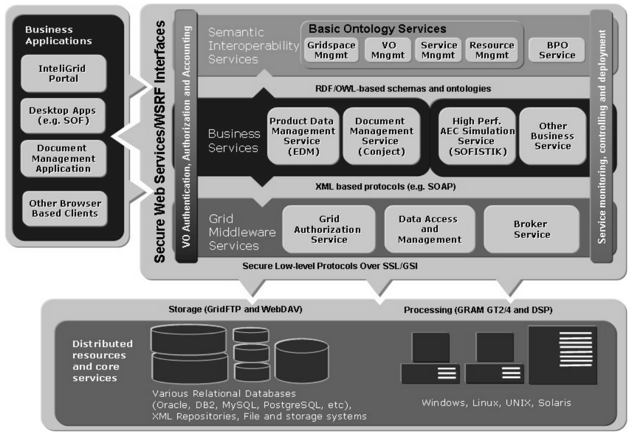

Figure 3. Ontology services in the InteliGrid infrastructure (Turk et al. 2006).

The Ontology Services provide generic and specific convenience methods to create, manipulate and manage ontology instances of classes defined in the described ontologies. Interfaces make use of the XML-based OWL notation for data exchange. The SPARQL query language (Prud'hommeaux & Seaborne 2007) that is somewhat comparable to SQL is used for ontology querying. Convenience methods provide access to instances of dedicated ontology classes and facilitate their management, e.g. *createProject(...)*, *deletePerson()*, etc. Such methods are important for the support of less powerful clients (without elaborate semantic functionality) and to enable acceptance of semantic services in general. However, to exploit the full potential of the Ontology Services, generic methods delegating more semantic logic to the clients have to be used.

From technical point of view, the Ontology Services follow the philosophy of Service Oriented Architectures (SOA), i.e. common access is provided through a Web Service Interface with grid security, making the Ontology Services pervasive and platform neutral.

Within all developed services, structured ontology management on model and entity level is realised via the *Jena Semantic Web Framework for Java* (Jena 2005). It provides a generic ontology graph for object-oriented and full-model access. Based on the ontology graph, a lightweight reasoner engine (part of Jena) enables the derivation of additional assertions entailed in the model. The primary use of this mechanism is to support the inference process of deriving additional facts from instance data and class descriptions.

Persistent storage for ontology data and models of the Ontology Services is provides by a back-end database (MySQL). Database integration is part of the Jena Framework realised by an abstract Java interface for model management on ontology databases.

For clients based on Java, a software package was developed that facilitates the handling of ontology instances based on the deployed ontologies. It enables bi-directional mapping between the XML-based OWL notation (used by the Service Interfaces) and Java objects. In contrast to the Jena methods, explicitly revealing the triple-based logic model in the background, the newly developed package allows manipulating instances of ontology classes and their properties in normal object-oriented manner. This approach does not support the full possible expressiveness of the developed ontologies but experience from the InteliGrid project clearly shows that the acceptance of semantic technology rises remarkably, once client developers can use the mapped object-oriented instance model.

For more details about the deployed Ontology Services, see (Gehre & Katranuschkov 2007).

4 APPLICATIONS UTILISING THE ONTOLOGY SERVICES

The developed ontologies and ontology services establish the conceptual and architectural backbone of a semantic infrastructure. They facilitate information management, improve the consistency of the distributed environment and make it less prone to errors. However, end-user applications can also strongly benefit from the added semantic value. The technology is well suited to support human-computer interactions because semantic models are more related to end-user perceptions than the usually applied, IT-biased database schemas.

To demonstrate this, a set of client applications were developed. They apply the ontology models directly (1) for management of the VO grid environment by administrators, and (2) for engineering tasks performed by end-users. A *Gridspace Management Client* supports the grid administrator in the process of managing users, resources and VO projects residing on the grid environment. A *VO Management Client* provides similar functionality but restricted to a particular VO, additionally providing role-based access control and permission management. At last, the developed *Business Process Object Manager* can be used to browse through and execute Business Process Objects in a structured way. These three client tools are described in more detail in the remainder of this chapter.

4.1 *Gridspace management client*

The Gridspace Management Client is dedicated to managing general semantic organisational information regarding the grid environment itself, i.e. information about the entities registered in the grid space. The grid space is the breeding environment in which new VOs are created. Main issues for that client are the management of users and services registered on the grid, as well as general management capabilities for all VO projects residing on the grid space (such as create/update/delete VO project). The two screenshots of the client (Figure 4 and Figure 5 below) illustrate the creation of a new VO and the browsing through the semantic information about registered resources.

As can be seen on Figure 5, using ontologies for the semantic description of resources in the grid allows advanced classification mechanisms, flexible sets of datatype and (referential) object properties, as well as presentation techniques that are easily understandable for end-users, even without much information transformation efforts for the end-user interfaces.

The client tool uses the Gridspace Management Service and the Resource Management Service in the background, i.e. semantic interoperability is realised on the basis of the *Organisational* and the *Resource Ontology*.

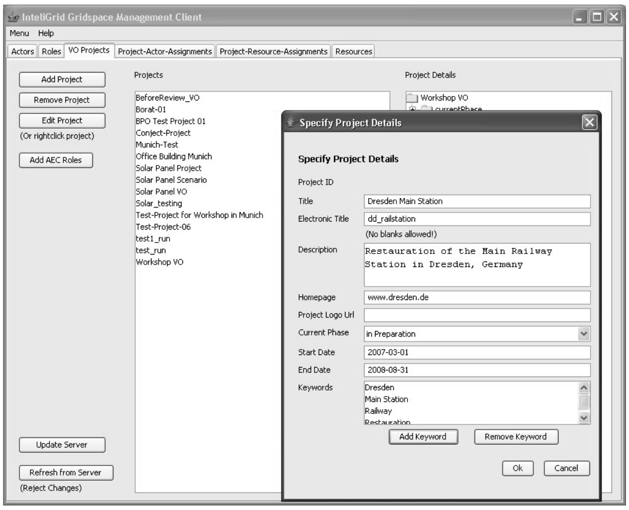

Figure 4. Screenshot of the Gridspace Management Client: Creating a new VO Project.

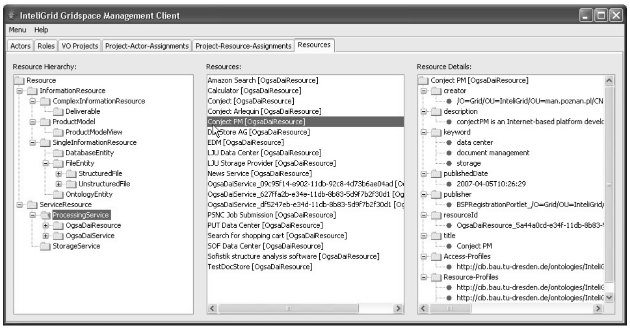

Figure 5. Screenshot of the Gridspace Management Client: Resources in the Gridspace.

4.2 *VO management client*

Particular VOs can be managed by the VO Management Client. Its focus lies on forming and managing the structure of the VO by assigning roles to actors, granting permissions to roles, etc. The actor-role-permission approach is implemented on the basis of the RBAC standard (Ferraiolo et al. 2001).

In contrast to the Gridspace Management Client, domain-specific extensions are more important for the VO Management Client because of the inherent organisational polymorphism of VO projects. Hence, the developed generic *Organisational Ontology* has to be carefully reviewed during the setup phase of a VO project and extended as necessary. If dedicated concepts are only specialised, further adaptation of service methods will not be necessary. The applied semantic web approach enables this largely. However, if specific organisational structures of a VO project diverge heavily from the predefined ontology structures, more efforts will be typically needed to integrate the required domain-specific features.

The VO Management Service uses information stored in the Grid Middleware Services in order to keep its data up to date and consistent with that stored in more fine-grained security services as e.g. the Grid Authorization Service or a Product Data Management Service.

The screenshot of the VO Management Client shown on Figure 6 below demonstrates the process of granting resource permissions by assigning resources to roles. All actors that possess this role are allowed to use the assigned resources afterwards.

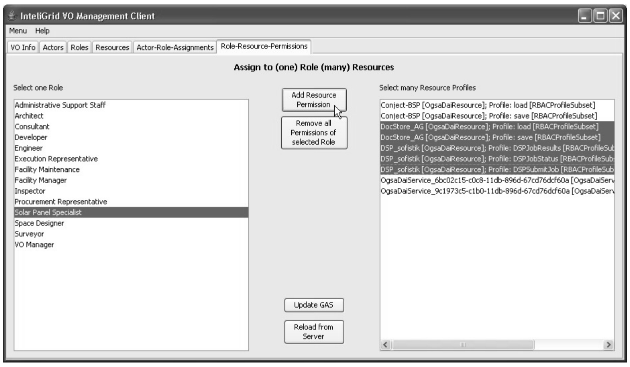

Figure 6. Screenshot of the VO Management Client: Granting Role-Resource Permissions.

As can be seen on the right side of the figure, more than one *RBAC Resource Profile* can be defined, i.e. by granting permission to a specific profile of a resource a specific role may get only read-access while other roles may get read/write access. More information about the developed flexible *Resource Profile* concept and the implemented set of pre-defined resource profiles as part of the *Resource Ontology* is provided in (Gehre et al. 2006).

4.3 *BPO management client*

The BPO Management Client provides for straightforward use of the BPO Services. Its GUI is divided into four main parts (see Figure 7 below). Panel 1 (left) provides a tree view of all instantiated user roles in the *Organisational Ontology* together with a high-level classification of their relevant business tasks. Panel 2 (top centre) lists all BPOs for the focused element in Panel 1 whereby the list automatically expands to reflect role inheritance. Panel 3 (top right) includes fields providing some details of the focused BPO and, more importantly, the elementary executable processes the BPO expands to. Thus, even on this, purely ontological level, end-users are provided valuable guidance for both their stand-alone work and the collaboration needs and demands. At last, Panel 4 (below 2 and 3) provides for actual BPO execution via either built-in or external plug-in services and tools, utilising the *Resource Ontology* and the *Service Ontology* specifications. The implementation of this functionality is provided via a generic *Factory object* that needs to be respectively applied for each BPO for which direct service execution from the client needs to be supported.

The sample BPO presented on Figure 7 shows the extraction of an IFC model view from a Product Model Server in the InteliGrid environment. The developed approach enables processing the BPO without asking the user to remember where the data and services are stored. This information is automatically provided by the related Resource Ontology Service that stores detailed access profiles for the resources described as ontology instances, i.e. the semantic system takes care of the place(s) where the resources are stored as well as of the data needed to access them. More information about the developed *Business Process Objects* concept is provided in (Katranuschkov et al. 2006) and in (Gehre et al. 2006).

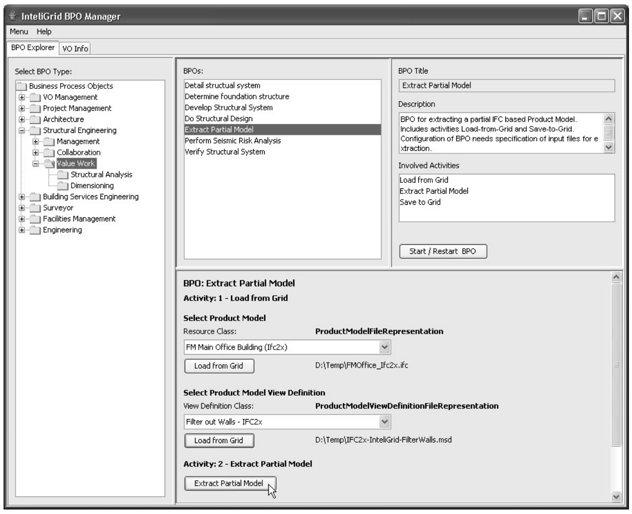

Figure 7. Screenshot of the BPO Manager: Browsing and execution of Business Process Objects.

5 CONCLUSIONS

The preceding chapters of this paper presented an overview of the authors' research on a semantic framework for VO management. Innovative aspects of the developed framework are: (1) the consistent definition and application of a set of newly defined semantic web ontologies for SOA-based grid environments for VO-centred business, (2) the design and development of Ontology Services with grid security features providing semantic information about entities in the grid (ontological metadata), organisational aspects forming the business structure of the Grid environment and residing VOs, and coherently integrated semantic Business Process Objects, and (3) client applications demonstrating the applicability and potential of the framework for semantic VO interoperability for a number of activities related to Gridspace and VO management as well as the management and execution of integrated, semantically defined Business Process Objects.

Design and prototype implementation convincingly provided proof of concept, but further development efforts are still needed regarding better support of domain-specific ontology extensions, definition and implementation of a larger set of Business Process Objects, more flexible semantic human-computer interaction, and finally, the study of adoption factors to achieve full-scale results.

ACKNOWLEDGEMENTS

This work has been partly funded by the European Commission through the FP6 project InteliGrid (IST-004664, 2004-2007). The support of the funding agency and of the partners in the InteliGrid project is herewith gratefully acknowledged.

REFERENCES

Berners-Lee T., Hendler J. & Lassila O. 2001. The Semantic Web. *Scientific American 284* (May/2001), 34-43.

Camarinha-Matos L. & Afsarmanesh H. 1999. The Virtual Enterprise Concept. In: *Infrastructures for Virtual Enterprises*. Kluwer Academia Publishers, ISBN 0-7923-8639-6.

Ferraiolo D.F., Sandhu R., Gavrila S., Kuhn D. R. & Chandramouli R. 2001. Proposed NIST Standard for Role-Based Access Control. *ACM Transactions on Information and System Security*, Vol. 4, No. 3, August 2001, pp. 224-274, Online available at: http://csrc.nist.gov/rbac/rbacSTD-ACM.pdf

Foster, I., Kesselman, C. & Tuecke, S. 2001. The Anatomy of the Grid: Enabling Scalable Virtual Organizations. *International Journal of High Performance Computing Applications*, 15 (3). 200-222.

Gehre A., Katranuschkov P., Stankovski V. & Scherer R. J. 2005. Towards Semantic Interoperability in Virtual Organisations. In: Scherer R.J, P. Katranuschkov & S.-E. Schapke (eds.) *Proceedings of CIB-W78 22nd Conference on Information Technology in Construction*, CIB Publication, ISBN 3-86005-478-3, pp. 307-314.

Gehre A., Katranuschkov P., Wix J. & Beetz J. 2006. *InteliGrid Deliverable D31: Ontology Specification.* The InteliGrid Consortium, c/o Univ. Ljubljana, Slovenia.

Gehre A. & Katranuschkov P. 2007. *InteliGrid Deliverable D32.2 – Ontology Services*. The InteliGrid Consortium, c/o University of Ljubljana, Slovenia.

Herman I. 2004. *Introduction to the Semantic Web*. W3C Presentation 27 Oct 2004. Online available at: http://www.w3.org/Consortium/Offices/Presentations/SemanticWeb/Overview.html

Jena. 2005. *Jena - A Semantic Web Framework for Java.* Online available at: http://jena.sourceforge.net/

Katranuschkov P., Gehre A., Keller M., Schapke S.-E. & Scherer R. J. 2006. Ontology-Based Reusable Process Patterns for Collaborative Work Environments in the Construction Industry. In: Cunningham P. & Cunningham M. (eds.) *Exploiting the Knowledge Economy*, IOS Press, ISBN 1-58603-682-3, pp. 1055-1063.

Martin D., Burstein M., Hobbs J., Lassila O., McDermott D., McIlraith S., Narayanan S., Paolucci M., Parsia B., Payne T., Sirin E., Srinivasan N. & Sycara K. 2005. *OWL-S: Semantic Markup for Web Services*. Online available at: http://www.daml.org/services/owl-s/1.1/overview/

Miller E. 2003. *Weaving Meaning: Semantic Web Applications*. Presentation at INTAP, November 11, 2003, Tokyo, Japan. Online available at: http://www.w3.org/2003/Talks/1117-semweb-intap/

Prud'hommeaux E. & Seaborne A. 2007. *SPARQL Query Language for RDF.* W3C Working Draft, 26 March 2007. Online available at: http://www.w3.org/TR/2007/WD-rdf-sparql-query-20070326/

Turk Z. (ed.) et al. 2006. *InteliGrid Deliverable D13.2 – Semantic Grid Architecture*. The InteliGrid Consortium, c/o University of Ljubljana, Slovenia.

W3C 2004a. *Web Ontology Language (OWL).* Online available at: http://www.w3.org/2004/OWL/ .

W3C 2004b. *RDF Vocabulary Description Language 1.0: RDF Schema.* W3C Recommendation 10 Feb 2004, Online available at: http://www.w3.org/TR/rdf-schema/