Semantic Grid – Convergence of Technologies

Dagstuhl Seminar 05271, 03.07. – 08.07.2005

Sponsored by EU IST network Knowledge Web

York Sure[1], Carole Goble[2], and Carl Kesselman[3]

[1]Universität Karlsruhe, D
sure@aifb.uni-karlsruhe.de

[2]University of Manchester, GB
carole@cs.man.ac.uk

[3]USC/Information Sciences Institute, USA
carl@isi.edu

1 Motivation & Overview

The scientific paradigms of the Semantic Web, Web Services, Agents, Peer-to-Peer Networks and Grid Computing are currently receiving a lot of attention in the research community, and are producing solutions to important problems ranging from e-science to e-business. The United States DAML program, the European Commission and other organisations have also been investing heavily in these technologies. This Dagstuhl Seminar brought together world-leading experts from the diverse organizations and research areas. It strengthened the international collaboration with the aim to realize the vision of the Semantic Grid.

In this section we briefly introduce the vision for a Semantic Grid followed by the main research topics which potentially converge into a Semantic Grid. We then give an overview on the agenda of the week in Section 2. We summarize the presentations in Section 3 and the break-out sessions in Section 4, followed by a summary of selected highlights in Section 5. Finally we conclude in Section 6. Attached is a list of participants.

1.1 The Vision: Semantic Grid

A Grid depends on understanding the resources it has available, their capabilities and how to assemble them. Thus Grid middleware thrives on metadata. Currently this metadata is largely managed in an ad hoc way by catalogues, registries or other forms of information services. This makes it hard to share, and interpret by services other than the originators.

Often these schemas are fixed, which makes them rather inflexible. Much of the metadata is hard-coded and buried in code libraries, type systems, or grid applications. This makes it hard to adapt and configure. Finally, understanding and know-how is frequently tacit, embedded in best practice and experience rather than explicitly recorded. This makes sharing and adaptation extremely difficult. Thus, existing Grid Services deal with knowledge in the form of metadata and its associated semantics in an implicit fashion, providing poor mechanisms for sharing this knowledge with other Grid components.

The Semantic Grid[1] is an initiative to systematically expose semantically rich information associated with resources to build more intelligent Grid services. It is an extension of the current Grid in which information and services are given well defined and explicitly represented meaning, better enabling computers and people to work in cooperation (see e.g. [6]).

Semantic Grids not only share computational and data resources, but also explicitly share and process metadata and knowledge. The Semantic Grid primarily aims to add meaning (ontologies, annotations and negotiation processes as studied in the Semantic Web and Software Agent paradigms) to the Grid. In this way, the Semantic Grid not only provides a general semantic-based computational network infrastructure, but a rich, seamless collection of intelligent, knowledge-based services for enabling the management and sharing of complex resources and reasoning mechanisms.

In the Semantic Grid knowledge and semantics are deployed explicitly for Grid applications and for the development of innovative Grid infrastructure. This knowledge-oriented semantics-based approach to the Grid goes hand-in-hand with the exploitation of techniques and methodologies from intelligent software agents representing various components of the virtual organizations and interacting in a P2P way.

In recognition of the potential importance of Semantics in Grids, the Global Grid Forum[2] standards body chartered a Semantic Grid Research Group in 2003 (GGF SEM-GRD RG)[3]. In the last few years, several projects have embraced this vision

[1] http://www.semanticgrid.org
[2] http://www.ggf.org/
[3] http://forge.gridforum.org/projects/sem-rg/

and there are already successful pioneering applications that combine the strengths of the Grid and of semantic technologies [10].

1.2 Semantic Web

"The Semantic Web aims to bring structure to the meaningful content of Web pages, creating an environment where software agents roaming from page to page can readily carry out sophisticated tasks for users" [1]. Key to the Semantic Web are "ontologies" [11]. A large number of projects (e.g. the EU integrated project SEKT[4], the UK AKT integrated Research Collaboration[5], and the US DAML initiative[6]) are aiming to solve problems like ontology alignment and mapping [4], reasoning over inconsistent models [7] or ontology learning [2], to name but a few.

1.3 Web Services

Web Services promise a new level of service on top of the current web. However, in order to employ their full potential, appropriate description means for web services need to be developed. Current technologies such as UDDI, WSDL, and SOAP provide limited support in mechanizing service discovery, service orchestration, service comparison and automated negotiation. Thus, current efforts (e.g. in the EU integrated project DIP[7]) aim to use the strengths of Semantic Web technologies to enable Semantic Web Services (see e.g. [9]).

1.4 Grid Computing

Grid Computing is a new field concentrating on "flexible, secure, coordinated resource sharing among dynamic collections of individuals, institutions, and resources - what we refer to as virtual organizations" (cf. [5]). The term "Grid" is used to mean many things. Here are some definitions from [6]: "The" Grid refers to the concept and the vision, and as such does not exist as an artifact. "A" Grid is a particular virtual organisation (VO) of heterogeneous distributed resources that will collaborate to solve a problem.

A VO of machines linked, by a high performance network, to form a virtual machine, such as the TeraGrid[8], is a Computational Grid. Other VOs reflect resources (a Data Grid forms a virtual database), geography (National Grid Service of UK resources), fields (Mouse Genome Grid), disciplines (BioGrid), or problems (protein folding simulation). There are many interoperating Grids. Production Grids are Grids of resources used routinely, commonly using previous generation middleware. They currently solve hetereogeneity and dynamic resource problems by enforcing rigorous conditions on participation, and using hardwired middleware.

[4] `http://www.sekt-project.com`
[5] `http://www.actors.org/akt`
[6] `http://www.daml.org`
[7] `http://dip.semanticweb.org/`
[8] `http://www.teragrid.org/`

Grid middleware infrastructure is the software services stack, policies, protocols, standards and APIs that make Grids. Reference implementations include Condor[9], Globus[10] and Unicore[11]. Grid tools, such as resource heartbeat monitors and portals, enable the management and use the Grid infrastructure. Finally, Grid applications use a Grid by means of its middleware to solve a problem, often a scientific one.

Initial research on Grid Computing focused on hiding the heterogeneity of computational resources and providing large scale data and computation systems. Current topics include support for a globally distributed collaboration, a service oriented approach and information management issues.

As a consequence, one of the key challenges in today's Grids is the need to deal with knowledge and data sources that are distributed, heterogeneous, and dynamic, and where effective elicitation of implicit knowledge is a necessary component of the overall system. In such systems, a complete global viewpoint or understanding is impossible to achieve. We therefore need to go beyond centralised knowledge service provision, and develop effective, open, distributed, knowledge-based solutions.

2 Agenda

	Monday *4th July 2005*	**Tuesday** *5th July 2005*	**Wednesday** *6th July 2005*	**Thursday** *7th July 2005*	**Friday** *8th July 2005*
07:30-08:45h	*Breakfast*				
09:00-10:30h	Welcome by Organizers *Speaker: York Sure* The Grid *Challenges and Opportunities* *Speaker: Carl Kesselman*	What can the SW do for the Grid? *Speaker: Jim Hendler* What can the Grid do for the SW? *Speaker: David De Roure*	Break-out Groups	(start at 9:30h) Break-out Groups	Final Wrap-up and Farewell *Chair: Carole Goble and York Sure*
10:30-10:45h	*Coffee*				
10:45-12:15h	Semantic Grid Roadmap *Speaker: Ziga Turk*	What Agents can do for Grids *Speaker: Yolanda Gil* Application Use Cases *Speaker: Jim Myers*	Break-out Groups	Break-out Groups	
12:15-14:00h	*Lunch*				
14:00-15:30h	Semantic Web Tutorial *Speaker: Jim Hendler* Grid Brief *Speaker: Andrew Grimshaw*	Middleware Use Cases *Speaker: Ian Foster* Where does semantics fit into NextGrid? *Speaker: David Snelling*	**Excursion**	Break-out Groups	Departure
15:30-16:00h	*Coffee*				
16:00-18:00h	Why Grids for Distr. Int. Systems: Because Robustness matters! *Speaker: Yolanda Gil* WSMO and the GRID *Speaker: John Domingue*	Where does semantics fit into the OGSA? *Speaker: Andrew Grimshaw* EU Frame Program 7 *Speaker: Eoghen O'Neill*		Break-out Groups	
18:00h	*Dinner*				
22:00h	*Nightly Cheese*				

Figure 1: Agenda of the Semantic Grid Dagstuhl Seminar

The agenda of the seminar was designed to quickly bring participants to a certain level of understanding and then leave them lots of time for in-depth discussions on specialized topics identified ad-hoc during the seminar (see Figure 1).

The seminar started with pre-selected introductory presentations And ad hoc tutorials covering the various communities and topics to "swing in" all participants and bring

[9] http://www.cs.wisc.edu/condor
[10] http://www.globus.org/
[11] http://www.unicore.org/

them on a first shared level of understanding each other. Then, a number of break-out sessions on most relevant topics were discussed in small groups. The sessions were suggested and selected during the seminar in a lively and interactive manner.

Approximately one third of the time was dedicated to presentations, one third to break-out groups and one third to the wrap-up session and the traditional social event on Wednesday.

3 Summary of Presentations

On Monday morning **York Sure** welcomed the participants on behalf of the organizers and opened the seminar. **Carl Kesselman** emphasized in his presentation "The Grid" the resource sharing aspect of grids as *highly dynamic resource sharing across multi-institutional boundaries*. **Ziga Turk** presented "Semantic Grid Roadmap" work performed in the project IntelliGrid and pointed out that roadmaps need to be *driven by needs, not by solutions*. In the following discussion, missing *lifecycle* aspects in the roadmap were pointed out which led to the instantiation of a break-out group having lifecycle aspects as topic. In addition to the planned agenda **Jim Hendler** gave a tutorial on "Semantic Web" followed by a tutorial given by **Andrew Grimshaw** called "Grid Brief". **Yolanda Gil** argued for robustness in her presentation "Why Grids for Distributed Intelligent Systems: Because Robustness matters". **John Domingue** gave insights on "Web Services Modelling Ontology (WSMO) and the Grid".

On Tuesday **Jim Hendler** bridged between Semantic Web and Grid in his presentation "What the Semantic Web can do for the Grid" and demonstrated as potentially useful Semantic Web features for Grids e.g. *RDF/S as very extensible metadata model* and *recursive SPARQL queries for navigational purposes*. **David De Roure** emphasized in his presentation "What the Grid can do for the Semantic Web" missing pieces in grids such as *negotiation, autonomy and self-organization* and included a brief history of the Semantic Grid. **Yolanda Gil** bridged between agent and grid communities in her presentation "What Agents can do for Grids: Cognitive Grids" and argued that in cognitive grids one usually adds *as much semantics as needed*. **Jim Myers** discussed use cases in his presentation "What can the Semantic Grid do for Science and Engineering", in particular the Scientific Annotation Middleware (SAM) at the National Center for Supercomputing Applications. **Ian Foster** presented further use cases in "Middleware Use Cases" where he explained that often there is *little or no a priori knowledge about the middleware environment* and metadata would be needed to describe *resource and service properties, data formats etc.*. Further he stressed that *discovery and negotiation* are key for successful applications. He mentioned that currently more and more data become available from monitoring Grid environments, but it is unclear what to do with it. **Eoghen O'Neill** showed a glimpse to the EU Frame Program 7 in his presentation "The Semantic Grid in the context of EU IST Research" where he also highlighted the funding situation in the upcoming 5th EU IST Call for Grid-related projects. **David Snelling** reported from the NextGrid project in his presentation entitled "Semantics and NextGRID". **Andrew Grimshaw**s presentation "Semantics, OGSA, and Agents (Oh my)" included both perspectives, the potential benefits of Semantics for Grids (e.g. *type conversions*) and the potential benefits of Grids for Semantics (e.g. using a

Grid as *scalable execution environment* for semantic technologies).

Further details can be found in presentations, abstracts and papers to be found at the materials page of the Semantic Grid Dagstuhl Seminar website `http://www.dagstuhl.de/05271/Materials`.

4 Summary of Break-out Sessions

During the seminar all participants decided ad-hoc on the topics for break-out sessions. The following sessions were agreed upon and held at the seminar.

- **Security and Trust**
  organized by Daniel Olmedilla, Omer Rama and Wolfgang Nejdl
- **Datagrids and Digital Libs**
  organized by Reagan Moore
- **Virtual Organisation Lifecycle**
  organized by Carl Kesselman
- **Scientific Workflows**
  organized by Bertram Ludaescher and Yike Guo
- **Marriage of OWL and Resource Properties**
  organized by David Snelling and Andrew Grimshaw

We would like to thank the organizers of the sessions for their efforts and their summaries of the sessions.

4.1 Security and Trust (inc. SAML)

Current trend in GSI is to enable trust relationships to be established in the Grid community – generally through the use of X509-based digital certificates, and more recently, through the use of security assertions (SAML) and role-based access management (PERMIS and Shibboleth). However, those security mechanisms still do not scale. Among the existing problems we can identify mechanisms that are too rigid for authentication and authorization, in terms of access control, and the lack of an ability to determine how "trustworthy" the result obtained from a specific provider is likely to be.

Trust management provides us with the basis to overcome these two points of view. However, the general notion of "trust" is excessively complex, and appears to have many different meanings depending on how it is used. Trust is seen as a multifaceted issue and may be related to other themes such as risk, competence, security, beliefs and perceptions, utility and benefit, and expertise. In addition, policy-based trust management is understood as statements guiding a process where two strangers are able to commit a specific transaction.

Therefore, the aim of this special session was to encourage the discussion about and identify the advantages/uses/requirements/threads of applying trust on Grid computing from the following two complementary points of view:

- **Access Control:** some of the existing problems can be addressed by extending Grid Security Infrastructure with trust negotiation mechanisms. These extensions, can provide the Grid with property-based authorization mechanisms, automatic gathering of required certificates, bidirectional and iterative security negotiations and policy based authorizations.

- **Provision of Service:** the notion of trust transcends beyond the restrictive security issues that are currently being explored in the Grid community. Where security issues are primarily concerned with ensuring that the result being provided is to come from a traceable source, trust issues can also relate to the degree of belief a user has in a particular provider, and is therefore much more subjective in nature. Some questions to be answered would be: "Can a trust rating be associated with a Grid Service? If so, how is this calculated and what does it mean?"

Active discussions including different points of view and backgrounds led our session to the following results:

- **A working definition of Trust**, which removes its ambiguity

- **A list of requirements** in order to address efficiently trust management issues in current Grids

- **An identification of the phases within a Virtual Organization lifecycle** where trust is a challenge

- **A roadmap of actions** to be undertaken aiming at solving those challenges

4.2 Datagrids and Digital Libs

The relationship between semantic web technology, the data management mechanisms supported by data grids, and the data discovery mechanisms supported by digital libraries was examined. An insight was developed into the characterization of grid-related systems:

- grids support virtualisation of workflows

- data grids support virtualisation of data

- semantic webs support virtualisation of information

- semantic grids support reasoning on inferred attributes

From the perspective of data grids, an additional level of virtualisation is needed, that of the state information maintained by the data grid to track operations performed upon data. While semantic grids support logical reasoning on logical relationships between labelled objects, data grids need to support causal reasoning on procedures or operations applied to labelled objects. The semantic grid technology can be applied to the management of state information about data and workflow, provided name spaces are developed to name operations, and logical relationships are defined between the operations.

A simple example of the type of reasoning that is desired is the concept of persistent objects. The preservation community desires the ability to apply current display applications to obsolete data formats. By characterizing both the structure of an object and the allowed operations on the structures, this can be achieved. A data display virtualisation can be achieved by mapping from the operations that a given display application wants to perform, to the allowed operations that can be performed upon a characterization of the structures in the object.

By naming the structures, semantic grids may be able to do the reasoning that asserts the allowed operations on the structures, which are then mapped to the operations that the display application needs to perform.

4.3 Virtual Organisation Lifecycle

This session was in two parts. Three main topics of discussion were: What is a VO, what is the life cycle of a VO and how might Semantic Web technologies assist in VO management.

In the first session a case study from the earth Sciences Grid focused the discussion. In the second the meeting adopted the Dagstuhl seminar itself as an example of a VO has would be helpful to complete as an exercise to illustrate the VO lifecycle. Discussion first concentrated on clarifying what a Virtual Organisation is, what the priorities of a VO are and what/who are members of a VO.

A virtual organization (VO) comprises a set of individuals and/or institutions having direct access to computers, software, data, and other resources for collaborative problem-solving or other purposes. VOs are a concept that supplies a context for operation of the Grid that can be used to associate users, their requests, and a set of resources. The sharing of resources in a VO 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.

Line Pouchet and Carl Kesselman gave an example of a VO for the Earth Systems Grid, which was bolted together by hand and took two years to form. Moreover, the criteria for membership were hard to elicit and consequently hard to encode in policy. The chief barrier was sociological, getting people to agree and building trust. Thus any computer-assisted VO forming mechanisms must encompass organisation and process flows as well as information flows.

Jim Hendler sketched a conceptual figure that showed that the Semantic Grid activity is effectively trying to explicitly model the VO, through reconciling or at least bridging multiple models of policy, resource and credentials to enable interoperability without full integration.

Jim Myers produced a list of things that are shared in VOs. A second topic was the use of Semantic Web technologies – OWL, PSL – to represent and reason over controlled vocabularies, ontologies (models) and processes to model a VO.

Line presented an OWL version of the Earth Sciences VO model using the Maryland SWOOP editor. Because of the well-defined formalism of OWL the audience were able to read off an unambiguous interpretation by sight-reading. This was considered a major benefit of adopting these technologies. Encoding credentials and policies in OWL and then reason about membership was thought to be a major benefit to the VO forming process.

Questions that reasoning would assist with include: Can you be a member of my VO given your credentials and my policy? How do I set your roles so you can be a member? Are these set of policies consistent or mutually inconsistent?

The third topic, addressed in the second session and using the Dagstuhl workshop as the motivating example, was the lifecycle of a VO. Specific processes representing the lifecycle could be encoded as workflows. Although OWL can declaratively describe the task and purpose of a workflow it is inappropriate to describe the flow itself; other languages like PSL or BPEL, or extensions such as OWL-SW (from the SIMDAT project) attempt this.

The conclusion is that VO membership and policy management is potentially one of the wins for Semantic Grids, and has a real need for the classification and rule based reasoning capabilities of the technologies. However, there are concerns whether policy can be reliably elicited to be encoded, and a suspicion that in many cases computer-assisted VO management is more practically feasible than automated management.

4.4 Scientific Workflows

Scientific workflows allow scientists to automate repetitive data management, analysis, and visualization tasks, and to document the provenance of analysis results. Scientific workflows are composed of interlinked computational components (sometimes called *actors*), and the datasets that are consumed and produced by those components. Scientific workflow systems are problem-solving environments to design, reuse, share, execute, monitor, and archive scientific workflows. As such, they are the primary tool that end user scientists use when interacting with the emerging e-Science cyberinfrastucture. Scientific workflow systems can often benefit from both, Grid and Semantic Web capabilities.

Thus, scientific workflows can bring together these otherwise loosely connected technologies and "catalyze the reaction" between them. For example, *compute-intensive* workflows (e.g. simulation models for quantum chemistry, protein folding, ecological niche modeling, weather forecasting, ocean currents, supernova explosions, etc.) require significant computational power available through Grid and/or dedicated cluster resources, and thus can benefit from Grid services for distributed and parallel execution and scheduling.

Scientific workflows can also be *data-intensive* (consuming and producing large amounts of data, thus suggesting the use of Data-Grid middleware), and *metadata-intensive*, i.e., requiring and producing information on the provenance of datasets and the meaning of data, in particular, to facilitate reproducability of experiments (e.g. by capturing essential parameters of experiment protocols) and for reuse of workflows and analysis products. Metadata-intensive workflows often employ Semantic Web standards such as RDF and OWL to capture metadata in a machine-processable form.

Taken together, resource management provided by Grid services and knowledge capture and management through Semantic Web technologies, provide essential capabilities of any general purpose, large-scale scientific workflow systems. In various ongoing projects we employ these technologies to enhance the KEPLER scientific workflow system and make it more versatile for the scientist and more interoperable with other e-Science/cyberinfrastructure tools and services.

4.5 Marriage of OWL and Resource properties

Goal of the session was to explore the potential for linking the current GRID architectural infrastructure as developed by partners of the GGF (Global Grid Forum) in a lightweight way with Semantic Web technologies. Of particular interest in the session was the usage of RDF/S and OWL to describe GRID meta data such as WSRF resource properties. Concrete efforts were started by modelling a lightweight ontology which is available at the Dagstuhl seminar website as '.owl' file. The ontology was an approach to bind ontological concepts from a resource ontology to resource properties.

The approach already has been adopted and enhanced and refined by the EU STREP project OntoGrid in deliverable D1.2 ([8]). Effectively, the resource ontology is seeding the GRID resource ontology activity that is led by EU STREP project UniGrid[12], particularly in efforts to combine the ontology developed with the CIM model [3].

5 Highlights of the Week

Figure 2: Group picture at the excursion

Highlights of the week included endless in-depth discussions at the atmospheric vine cellar, bicycle excursions to the nearby forests, the traditional "Wednesday-excursion" which consisted of a visit of the Villeroy & Boch museum, a boat-trip on the river Saar, a stop-over at the beautiful Saarschleife (see Figure 2) and a wine-tasting, and finally the marvellous "Thursday-barbecue" with local draught beer which culminated in the famous "lounge party" featuring lively dancing, art performance, and ontology engineering. The lounge party truly showed the convergence of communities as seminal

[12] http://www.unigrids.org/ontology.html

work on an ontology about Virtual Organisations was jointly performed by Grid and Semantic Web people.

By the way, if you are interested in “what are the roads not taken” by (most of the) participants of the seminar, we encourage you to have a look at the guestbook entries for the Semantic Grid Seminar during your next stay in Dagstuhl.

6 Conclusion

The main achievements of the Seminar include:

- creation of a strong and vivid Semantic Grid community, which shares understanding of principle ideas
- foundations for a WSRF-resource ontology
- foundations for a Virtual Organisation (VO) ontology, based on the EarthSciences Grid ontology
- an shared understanding of what a Semantic Grid is, and a need to encapsulate this in a bumper sticker
- an understanding of what a VO is, and using Dagstuhl as an example,
- the role of agents and Grids,
- a coordinated exchange of staff between Inteligrid, K-WfGrid and OntoGrid, which has already begun
- a book (working on it ...), and
- a number of visits stemming from the spin of the seminar (Frank Sieberlist visited Hannover and Manchester from Globus, for example)

The results of the workshop also contributed to presentations at the Knowledge Web Summer School on the Semantic Web 2005 on 22nd July, and the OECD Grid Global Science Forum in Sydney, Australia 25th September, both given by Carole Goble. Material developed in the VO discussions have contributed to a case study in the OntoGrid EU STREP and the WSRF/VO ontology has contributed to OntoGrid's Reference Architecture for Semantic Grids. Two EU IP proposals (BIG and BREIN) on the Semantic Grid have also be submitted in Sept 2005 by partners attending the workshop, based in part on discussions therein.

Contact details of all participants as well as contributions such as abstracts, presentations and further links can be found on the web site of the seminar:

`http://www.dagstuhl.de/05271/`

Acknowledgements: We thank the EU IST Network of Excellence Knowledge Web (see `http://knowledgeweb.semanticweb.org/`) for sponsoring our Semantic Grid Dagstuhl Seminar. We also thank the people at Schloss Dagstuhl for keeping our work focused by providing very pleasant surroundings and perfect service.

References

[1] T. Berners-Lee, J. Hendler, and O. Lassila. The semantic web. *Scientific American*, 2001(5), 2001. available at http://www.sciam.com/2001/0501issue/0501berners-lee.html.

[2] Paul Buitelaar, Philipp Cimiano, and Bernardo Magnini, editors. *Ontology Learning from Text: Methods, Evaluation and Applications*. Frontiers in Artificial Intelligence. IOS Press, 2005.

[3] Distributed management task force - common information model (CIM). Website, 2005. `http://www.dmtf.org/standards/cim/`.

[4] M. Ehrig and Y. Sure. Ontology mapping - an integrated approach. In C. Bussler, J. Davis, D. Fensel, and R. Studer, editors, *Proceedings of the First European Semantic Web Symposium*, volume 3053 of *LNCS*, pages 76–91, Heraklion, Greece, MAY 2004. Springer Verlag.

[5] I. Foster, C. Kesselman, and S. Tuecke. The anatomy of the Grid: Enabling scalable virtual organizations. *International J. Supercomputer Applications*, 15(3), 2001.

[6] C. Goble and D. De Roure. The semantic grid: Building bridges and busting myths. In *16th European conference on Artificial Intelligence (ECAI) 2004 including Prestigious Applicants of Intelligent Systems*, Valencia, Spain, 22-27 August 2004, 2004. PAIS IOS Press.

[7] P. Haase, F. van Harmelen, Z. Huang, H. Stuckenschmidt, and Y. Sure. A framework for handling inconsistency in changing ontologies. In *Proceedings of the Fourth International Semantic Web Conference (ISWC2005)*, LNCS. Springer, NOV 2005.

[8] I. Kotsiopoulos, S. Bechhofer, P. Alper, P. Missier, O. Corcho, D. Kuo, and C. Goble. Specification of a semantic grid architecture. OntoGrid Deliverable D1.2, University of Manchester, 2005.

[9] S. A. McIlraith, T. C. Son, and H. Zeng. Semantic web services. *IEEE Intelligent Systems*, 16(2):46–53, 2001.

[10] D. De Roure, Y. Gil, and J. Hendler. E-science special issue. *IEEE Intelligent Systems*, 19(1), 2004.

[11] S. Staab and R. Studer, editors. *Handbook on Ontologies in Information Systems*. International Handbooks on Information Systems. Springer, 2004.

Participants List

1. Jürgen Angele (ontoprise GmbH, D)

2. Bruce R. Barkstrom (NASA Langley Research Center, USA)
3. Sean Bechhofer (The University of Manchester, GB)
4. Marian Bubak (AGH – University of Science & Technology, PL)
5. Mario Cannataro (University Magna Graecia of Catanzaro, I)
6. Tim Clark (Harvard Medical School, USA)
7. Oscar Corcho (The University of Manchester, GB)
8. David De Roure (University of Southampton, GB)
9. John Domingue (The Open University, GB)
10. Alvaro A. A. Fernandes (The University of Manchester, GB)
11. Ian Foster (Argonne National Laboratory, USA)
12. Yolanda Gil (University of Southern California, USA)
13. Carole Goble (The University of Manchester, GB)
14. Asuncion Gomez-Perez (Universidad Politecnica de Madrid, E)
15. Andrew Grimshaw (University of Virginia, USA)
16. Yi-ke Guo (Imperial College London, GB)
17. Peter Haase (Universität Karlsruhe, D)
18. Heiko Haller (Universität Karlsruhe, D)
19. James A. Hendler (University of Maryland at College Park, USA)
20. Ian Horrocks (The University of Manchester, GB)
21. Kashif Iqbal (National University of Ireland, IRL)
22. Carl Kesselman (University of Southern California, USA)
23. Bertram Ludaescher (University of California, USA)
24. Jon MacLaren (Louisiana State University, USA)
25. Brian Matthews (Rutherford Appleton Lab., GB)
26. Reagan Moore (University of California at San Diego, USA)
27. Luc Moreau (University of Southampton, GB)
28. James D. Myers (Universtity of Illinois at Urbana Champaign, USA)
29. Wolfgang Nejdl (Universität Hannover and L3S, D)

30. Eoghan O'Neill (European Commission INFSO F2, B)
31. Daniel Olmedilla (Universität Hannover and L3S, D)
32. Axel Polleres (Universität Innsbruck, A)
33. Line Pouchard (Oak Ridge National Laboratory, USA)
34. Omer F. Rana (Cardiff University, GB)
35. David Snelling (Fujitsu Labs of Europe, GB)
36. Steffen Staab (Universität Koblenz-Landau, D)
37. Heiner Stuckenschmidt (Vrije Universiteit Amsterdam, NL)
38. Rudi Studer (Universität Karlsruhe, D)
39. York Sure (Universität Karlsruhe, D)
40. Domenico Talia (Universita della Calabria, I)
41. Valentina Tamma (University of Liverpool, GB)
42. Ziga Turk (University of Ljubljana, SLO)