

E-INFRASTRUCTURE: EUROPE MEETS THE E-SCIENCE CHALLENGE

GUEST EDITORS TONY HEY AND ANNE TREFETHEN

FEATURE ARTICLES

2 **The e-Science Challenge: Creating a Reusable e-Infrastructure for Collaborative Multidisciplinary Science**

Tony Hey, Corporate Vice President for Technical Computing, Microsoft Corporation
Anne Trefethen, Director, UK e-Science Core Programme, EPSRC

7 **OntoGrid: A Semantic Grid Reference Architecture**

Carole Goble and Sean Bechhofer, University of Manchester, UK

12 **How to Build an International Grid: Infrastructure, Applications & Community**

Fabrizio Gagliardi, EGEE Project Director, CERN
Bob Jones, EGEE Technical Director, CERN
Owen Appleton, EGEE Communications Officer, CERN

16 **SIMDAT**

Mike Boniface and Colin Upstill, University of Southampton IT Innovation Centre

21 **The Next Generation Grid**

Mark Parsons, NextGRID Project Chairman

NOTES AND COMMENTARY

26 **The Next Decade in HPC**

Craig Mundie, CTO, Microsoft Corporation

30 **PITAC's Look at Computational Science**

Dan Reed, Renaissance Computing Institute, University of North Carolina at Chapel Hill



Introduction

This issue of *CTWatch Quarterly* is intended to give an overview of the activity on e-Science and Grids in Europe. The European Commission were very early to identify Grids as a key technology for collaboration and resource sharing. The pioneering European DataGrid project, led by Fabrizio Gagliardi at CERN, worked with the world particle physics community, and with US computer scientists Ian Foster, Carl Kesselman and Miron Livny, to develop a global Grid infrastructure capable of moving large amounts of data and providing the vast shared compute resources needed to analyse this data. Many Petabytes of data per year will be generated by experiments at the Large Hadron Collider, currently under construction at the CERN Laboratory in Geneva. Similarly, the UK was also early to see the potential of Grid technologies for building the scientific 'Virtual Organizations' needed for networked scientific collaborations. In 2001, the UK announced the beginning of a \$400M 'e-Science' initiative - where the term e-Science was introduced by John Taylor, the Director of the UK's Office of Science and Technology, as a short-hand for the set of collaborative technologies needed to support the distributed multi-disciplinary science and engineering projects of the future.

It is now 2005 and the European Union has invested in a new generation of projects, building on the lessons of the European DataGrid and other similar projects. Besides investing in further R&D projects, the Commission has identified the need to develop and sustain an 'e-Infrastructure' consisting of a pan-European, high-speed research network, GEANT-2, together with a set of core Grid middleware services to support distributed scientific collaborations. The set of reports included here therefore includes three new R&D projects - SIMDAT, NextGrid and OntoGrid - plus the major Research Infrastructure project, EGEE - Enabling Grids for E-Science - which in many ways can be seen as the direct successor of the original European DataGrid project.

The SIMDAT project is developing generic Grid technology for the solution of complex application problems, and demonstrating this technology in several representative industry sectors. Special attention is being paid to security and the objective is to accelerate the uptake of Grid technologies in industry and services. Major European companies from the aerospace and automotive sectors and from the pharmaceutical industry are partners in the project which also involves major European Meteorological centers. By contrast, the NextGrid project is looking further out at the next generation of Grid technologies and is focused on inter-enterprise computing in the business sector with partners such as SAP, BT, Fujitsu, NEC and Microsoft.

OntoGrid represents another strand of activity and builds on pioneering work towards the development of a truly 'Semantic Grid' in the UK e-Science program. The project brings together knowledge services - such as ontology services, metadata stores and reasoning engines - with Grid services - such as workflow management, Virtual Organisation formation, debugging, resources brokering and data integration. This semantics-based approach to the Grid goes hand-in-hand with the exploitation of techniques from intelligent software agents for negotiation and coordination and peer-to-peer (P2P) computing for distributed discovery. These four projects are by no means all of the current EU Grid projects and details of these and other projects may be found on the Cordis website (<http://www.cordis.lu/>).

Tony Hey

Corporate Vice President for Technical Computing, Microsoft Corporation

Anne Trefethen

Director, UK e-Science Core Programme, EPSRC

The UK e-Science Program has now entered its third phase and this is focusing on laying the foundations for a sustainable national e-Infrastructure – or Cyberinfrastructure in US-speak. These activities are described in a short article by the editors.

Complementing the other article's in this issue, Dan Reed's personal reflections on the recent report of his PITAC subcommittee on Computational Science shows that a shared sense of current challenges and current opportunities is driving the development of e-Infrastructure and e-Science on both sides of the Atlantic.

The last article in this issue is of a different character and is a personal account by Microsoft's Chief Technology Officer, Craig Mundie, of his own roots in High Performance Computing and the reasons why Microsoft are now taking steps to become engaged with the HPC community. This is a fascinating glimpse into the future and indicates that Microsoft intends to play a major role in the development of 'commodity HPC' systems.

Tony Hey and Anne Trefethen

The e-Science Challenge: Creating a Reusable e-Infrastructure for Collaborative Multidisciplinary Science

1. Introduction

This issue of *CTWatch Quarterly* contains four articles that provide an overview of some of the major Grid projects in Europe. All these projects are aimed at developing distributed collaborative research capabilities for the scientists that are built on the deployment of a persistent middleware infrastructure on top of the high bandwidth research networks. The combination of a set of middleware services running on top of high speed networks is called 'e-Infrastructure' in Europe and 'Cyberinfrastructure' in the USA. In this brief article we shall abstract the key elements of such an e-Infrastructure from these projects and from our experience in our UK e-Science program. We look at the problems of creating and implementing a sustainable, global e-Infrastructure that will enable multidisciplinary and collaborative research across a wide range of disciplines and communities.

2. Background

The UK e-Science Initiative began in April 2001 and over the last four years, more than £250M has been invested in science applications and middleware development. In addition, the program created a pipeline from the science base to genuine industrial applications of this technology and, most importantly, has enabled the creation of a vibrant, multidisciplinary, e-Science community. This community comes together in its totality at the UK's annual e-Science All Hands Meeting which is held each September. We now have a community of over 650 who attend and join in to share experience and technologies. These meetings have brought together an exciting mix of scientists, computer scientists, IT professionals, industrial collaborators and, more recently, social scientists and researchers in the arts and humanities. Research scientists from all domains of science and engineering—particle physics, astronomy, chemistry, physics, all flavours of engineering, environmental science, bioinformatics, medical informatics and social science—as well as the arts and humanities are beginning to appreciate the need for e-Science technologies that will allow them to make progress with the next generation of research problems. In most cases, researchers are now finding themselves faced with an increasingly difficult burden of both managing and storing vast amounts of data as well as analyzing, combining and mining the data to extract useful information and knowledge. Often this can involve automation of the task of annotating the data with relevant metadata as well as constructing complex search engines and workflows that capture complex usage patterns of distributed data and compute resources. Most of these problems and the tools and techniques to tackle them are similar across many different types of application. It makes no sense for each community to develop these basic tools in isolation. We need to identify and capture a set of generic middleware services and deploy them on top of the high-bandwidth research networks to constitute a reusable e-Infrastructure. In the UK e-Science Initiative, this task—of identifying and implementing the key features of a national e-Infrastructure—was the remit of the Core Program.

The phrase e-Infrastructure—or Cyberinfrastructure in the US—is used to emphasize that these applications will be facilitated by a set of services that permit easy but controlled access to the traditional infrastructure of science—supercomputers, high performance clusters, networks, databases and experimental facilities. The e-Science challenge is to provide a set of Grid middleware services that are sufficiently robust, powerful and easy to use that application scientists are freed from re-inventing such low-level 'plumbing' and can concentrate on their

Tony Hey

Corporate Vice President for Technical Computing, Microsoft Corporation

Anne Trefethen

Director, UK e-Science Core Programme, EPSRC

science. A second challenge is to make this combination of middleware and hardware into a truly sustainable e-Infrastructure in much the same way as we take for granted the research networks of today.

3. Requirements for a Sustainable e-Infrastructure

The Grid projects referred to in these articles as well as the national e-Science programmes in Europe give us a good idea of what is required to create such a persistent, global, e-Science e-Infrastructure. In the UK the key elements have been identified as:¹

1. A competitive network of National Research and Education Networks (NRENs) together with their 'CERT' teams for security monitoring and emergency response. In the UK, the SuperJANET5 network and the CERT are run by UKERNA. Across Europe, the EU has funded the Dante organization to manage the GEANT2 network that connects the European NRENs.
2. A secure national and internationally accepted framework for multiple levels of authentication and authorisation. This must support both access within individual institutions as well as dynamic, cross-boundary 'Virtual Organizations' of research groups from different institutions.
3. A collection of software centres and repositories for open source reference implementations of open standards compliant, infrastructure middleware. This will require the participation or creation of organizations with a serious software engineering capability to research, support and maintain this middleware.
4. A national focus on digital 'curation' to provide scientists with support and guidance into the long-term preservation of both research data and traditional publications. By curation we mean annotation of data with metadata to enable efficient searching and provenance tracking.
5. Integrated access to national data sets and publications is emerging via the developing Open Access Subject and Institutional Repositories. Examples in the UK include the Arts and Humanities Data Service (AHDS), the Economic and Social Data Service (ESDS), the EDINA and MIMAS JISC funded services that offer sets of national data resources for education and research, and the resources of the British Library.
6. Remote access to large scale facilities such as Diamond and ISIS in the UK and the LHC, VLT and ITER internationally. Increasingly, scientists will have to pool their financial resources and perform experiments on facilities procured at a global level. For example, the particle physicists intend to use middleware developed in the EGEE project to create an LHC Grid for distribution and analysis from the machine in Geneva.
7. A set of national services both for HEC and Grid computing and for data services and long-term data archiving. High end supercomputers are clearly an important resource for computational scientists but there is also a need for more modest cluster resources.
8. National and international centres to enhance the creation of a strong culture of multi-disciplinary research and provide training in new informatics technologies. Much of the

¹ 'A National e-Infrastructure for Research and Innovation', Neil Geddes, Tony Hey, Anne Trefethen, Malcolm Read and Alan Robiette, *Discussion Paper for UK e-Science Steering Committee* [2004]

The e-Science Challenge: Creating a Reusable e-Infrastructure for Collaborative Multidisciplinary Science

new e-Science will be international and there is a need for a strong program of activity dedicated to building and educating a new multidisciplinary community of scientists.

9. Strong involvement in international standards activities both for infrastructure and for each of the global research communities. The GGF is focussing on developing a set of standards for infrastructure services while community organizations such as the International Virtual Observatory Alliance are delivering interoperability standards for their astronomy community.
10. Development of tools and services to support multidisciplinary and collaborative environments. These include portals providing access to quality data and services, national service and ontology registries and tools to support workflow and track provenance.

4. An Example: The Emerging UK e-Infrastructure

With funding from the e-Science Core Program, the UK is now in the process of implementing a prototype national e-Infrastructure. Several key components have been developed and these include:

4.1 An Open Middleware Infrastructure Institute

The e-Science Core Program has recently funded the continuation of an Open Middleware Infrastructure Institute (OMII-UK). In its second phase, the Institute is now built on three existing centres and leverages their joint user groups and the different competencies of the three teams. The lead partner is the original OMII at the University of Southampton which was set up in 2004 to provide well-engineered e-Science middleware sourced from the e-Science community.² They have now been joined by the OGSA-DAI team in Edinburgh that has developed middleware to support data access and integration now used worldwide;³ and by the myGrid project, which since 2001 has developed a set of workflow-based tools that have been widely adopted to support researchers in the bioinformatics community.⁴

² OMII: <http://www.omii.ac.uk/>

³ The OGSA-DAI project: <http://www.ogsadai.org.uk/>

⁴ The myGrid project: <http://www.mygrid.org.uk/>

By combining the expertise of these groups in OMII-UK, the e-Science Core Program has established a powerful source of well-engineered software, which should enable an integrated approach to the provision of higher level and more advanced tools. A dialogue is taking place with similar organizations in different countries, such as the NMI in the US and a new organization, OMII-China in Beijing.

4.2 A National Grid service

Established in April 2004, the National Grid Service (NGS) builds on the experiences of the UK e-Science community.⁵ At the core of the service there are two compute and two data clusters located at the Universities of Manchester, Oxford, and Leeds, and at the Rutherford Appleton Laboratory. These are supported by the Grid Operations Support Centre (GOSC) who maintain the UK e-Science Certificate Authority, a help desk, and provide training for administrators and IT professionals.

⁵ The National Grid Service: <http://www.ngs.ac.uk/>

The NGS service has now connected compute resources at several partner sites. Presently there are three such associate sites, namely; Cardiff, Lancaster, and Bristol Universities. In addition to this production service, other UK e-Science Centres play an important role in evaluating and testing Grid middleware as part of the software appraisal process for the NGS

and OMII-UK. Since the NGS has been in production mode, the number of registered users has risen to over 300 in a broad range of application areas.

At present, the core middleware of the NGS production grid is based on the Globus GT2 Toolkit and the SDSC Storage Resource Broker (SRB). As the Web Services versions of Grid middleware mature, it is expected that the NGS will migrate to a set of middleware services compliant with the GGF OGSA architecture. It is also expected that this set will include software from the OMII-UK as well as from the NMI and the EGEE project described in this publication.



4.3 A Digital Curation Centre

The Digital Curation Centre (DCC) has been established in Edinburgh by the e-Science Core Program and the JISC. Its role is to support best practice and to pursue research in data curation and digital preservation.⁶ In particular, it is working with different application communities to understand their specific challenges and identify best practice. The Centre will provide advice and support services to UK researchers and institutions. In the next five years, it is clear that many scientists are likely to be swamped with data. Managing the whole data chain, from acquisition and annotation through to integration and preservation, will be a major challenge. Tools to support collaborative working, workflow, provenance and high performance visualization will be needed. In some communities, there are business or legal requirements for long-term data preservation and access, as for example, with engineering drawings and clinical records.

⁶ The Digital Curation Centre:
<http://www.dcc.ac.uk/>

5. Conclusions: Embedding e-Science

At present, the e-Science research agenda both in technology and applications is largely being driven by leading-edge scientists and researchers who are prepared to engage with immature, 'bleeding edge' software and technologies. To engage a broader spectrum of the scientific community requires that the steepness of the learning curve be much reduced and the e-Science tools and technologies integrated into well known, familiar environments. Supportive, collaborative, 'virtual organizations' must be easy to establish and provide an adequate level of security and an acceptable user interface. Only with stable and robust middleware services will scientists be able to routinely construct the types of Grid that they need for their type of research.

Several other activities are underway in the UK that are attempting to move forward in this agenda of embedding e-Science into the fabric of research. These include:

1. A research and development programme in security for e-Science infrastructures and applications. Issues include GSI style digital certificates for VO membership to Shibboleth mediated trust networks between institutions.
2. A programme of research into usability issues related to tools, applications, e-Infrastructure and general methodologies.

The e-Science Challenge: Creating a Reusable e-Infrastructure for Collaborative Multidisciplinary Science

3. A program to develop flexible, easy-to-use Virtual Research Environments (VRE). The goal is to lower the barrier for adoption of the new e-Infrastructure services in several domains using portals to provide transparent access to resources.
4. Teaching and training courses to educate the next generation of e-Science researchers. Several universities now have Masters level programmes or components within such programmes that address some of the issues in e-Science. The National e-Science Centre (NeSC) also provides training for application scientists in new technologies as they emerge.⁷

⁷The National e-Science Centre:
<http://www.nesc.ac.uk/>

Similarly, there are many other EU R&D projects addressing a similar set of issues as well as a set of other national e-Science programs.

Given the large investment that the UK has made in e-Science since 2001, we are now beginning to see real benefits emerging for some application communities. This is true for projects both in the UK and the rest of the EU. Although some other application communities are still at an early stage of exploration of e-Science technologies, already the potential benefits are becoming clear for their particular area of research. The use of these technologies will have a profound change on the methodology and processes that the researchers have traditionally employed to do their science. With the advent of very large data sets, we are seeing a new form of data-centric, collections-based science begin to emerge to complement the traditional experimental, theoretical and computational approaches. There will be as much a change in social behaviour as a change in technology.

OntoGrid: A Semantic Grid Reference Architecture

What is a Semantic Grid?

The Grid aims to support secure, flexible and coordinated resource sharing by providing a middleware platform for advanced distributed computing. Consequently, the Grid's infrastructural machinery aims to allow collections of any kind of resources—computing, storage, data sets, digital libraries, scientific instruments, people, etc—to easily form Virtual Organisations that cross organisational boundaries in order to work together to solve a problem. A Grid depends on understanding the available resources, their capabilities, how to assemble them and how to best exploit them. Thus Grid middleware and the Grid applications they support thrive on the metadata that describes resources in all their forms, the Virtual Organisations, the policies that drive them, and so on, together with the knowledge to apply that metadata intelligently.

The Semantic Grid is a recent initiative to systematically expose semantically rich information associated with Grid resources to build more intelligent Grid services.¹ The idea is to make structured semantic descriptions real and visible first class citizens with an associated identity and behaviour. We can then define mechanisms for their creation and management as well as protocols for their processing, exchange and customisation. We can separate these issues from both the languages used to encode the descriptions (from natural language text right through to logical-based assertions) and the structure and content of the descriptions themselves, which may vary from application to application.

In practice, work on Semantic Grids has primarily meant introducing technologies from the Semantic Web² to the Grid. The background knowledge and vocabulary of a domain can be captured in *ontologies* – machine processable models of concepts, their interrelationships and their constraints; for example a model of a Virtual Organisation.³ Metadata labels Grid resources and entities with concepts, for example describing a job submission in terms of memory requirements and quality of service or a data file in terms of its logical contents. Rules and classification-based automatic inference mechanisms generate new metadata based on logical reasoning, for example describing the rules for membership of a Virtual Organisation and reasoning that a potential member's credentials are satisfactory.

In recognition of the potential importance of Semantics in Grids, the Global Grid Forum standards body chartered a Semantic Grid Research Group in 2003.⁴ The Forum's XML-based description languages such as the Job Submission Description Language, the Data Format Description Language and Oasis' Security Assertion Markup Language all identify the role of semantics. Their recent Database Access and Integration Services Working Group specification identifies the importance of semantics in integration, metadata management and discovery. In July 2005 the Grid and Semantic Web Communities came together in a week long Schloss Dagstuhl seminar.⁵

In the last few years, several projects have embraced the Semantic Grid vision and pioneered applications combining the strengths of the Grid and of semantic technologies, particularly the use of ontologies for describing Grid resources and improving interoperability.⁶ The UK myGrid⁷ project uses ontologies to describe and select web-based services used in the Life Sciences; the UK Geodise project uses ontologies to guide aeronautical engineers to select and configure Matlab

Carole Goble and
Sean Bechhofer
University of Manchester, UK

¹ C.A. Goble, D. De Roure, N. R. Shadbolt and A.A. Fernandes "Enhancing Services and Applications with Knowledge and Semantics" in *The Grid 2 Blueprint for a New Computing Infrastructure Second Edition* eds. Ian Foster and Carl Kesselman, 2003, Morgan Kaufman, November 2003.

² J. Hendler "Science and the Semantic Web," *Science* 299: 520-521, 2003.

³ L. Pouchard, L. Cinquini, B. Drach, et al., "An Ontology for Scientific Information in a Grid Environment: the Earth System Grid," *CCGrid 2003* (Symposium on Cluster Computing and the Grid), Tokyo, Japan, May 12-15, 2003.

⁴ <http://www.semanticgrid.org/>

⁵ <http://www.dagstuhl.de/05271/>

⁶ D. De Roure, Y. Gil, J. A. Hendler "Guest Editors' Introduction: E-Science," *IEEE Intelligent Systems* 19(1), Jan/Feb 2004: 24-25.

⁷ C. Wroe, C.A. Goble, M. Greenwood, P. Lord, S. Miles, J. Papay, T. Payne, L. Moreau. "Automating Experiments Using Semantic Data on a Bioinformatics Grid" *IEEE Intelligent Systems special issue on e-Science* Jan/Feb 2004.

OntoGrid: A Semantic Grid Reference Architecture

scripts;⁸ the Collaboratory for Multi-scale Chemical Science⁹ and CombeChem¹⁰ projects both use semantic web technologies to describe provenance metadata for chemistry experiments; the US-based Biomedical Informatics Research Network uses technologies to mediate between different databases in neuroscience;¹¹ and the UK's CoAKTing project uses ontologies to assist in virtual meetings between scientists.¹² On the Semantic Grid road we are now moving from a phase of exploratory experimentation to one of systematic investigation, architectural design and content acquisition for a semantic infrastructure that accompanies a cyberinfrastructure (Figure 1).

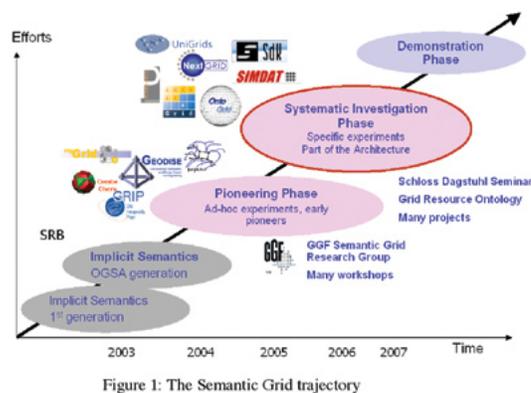


Figure 1: The Semantic Grid trajectory

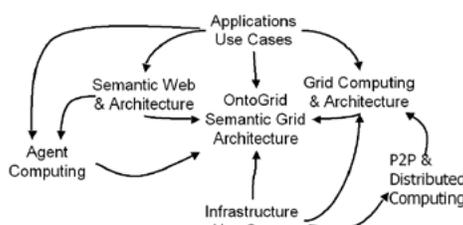


Figure 2: The influences on the OntoGrid project

OntoGrid

OntoGrid¹³ is an eight-partner EU FP6 project launched in October 2004 to investigate fundamental issues in Semantic Grids, bridging between the knowledge-based systems community and the Grid community. The project aims to show how knowledge technologies help deliver the next generation of Semantic Grid Computing systems and to experiment with the technological infrastructure needed for the development of knowledge-intensive, distributed open services for the Semantic Grid. The Semantic Grid should not only provide a general semantic-based computational infrastructure, but also a rich collection of knowledge services and knowledge-based services. Thus OntoGrid systematically brings together knowledge services (like ontology services, metadata stores and reasoning engines) with Grid services (such as workflow management, Virtual Organisation formation, debugging, resources brokering and data integration) adapted to semantic descriptions when they are available. This semantics-based approach to the Grid goes hand-in-hand with the exploitation of techniques from intelligent software agents and peer-to-peer (P2P) computing. OntoGrid mixes in techniques from agent computing for negotiation and coordination and peer-to-peer for distributed discovery (Figure 2).

OntoGrid is paving the way to Semantic Grids by investigating questions such as: Are semantic web technologies scalable? What's the impact of a semantic approach to legacy grids? How do we minimize the impact? What are the minimum knowledge services needed? What should be their capabilities? How do we harvest and tend the semantic content? Is there content that is common for all Grids and how much is application specific? How, when and where does a semantic approach add value to a "traditional" Grid approach? What is an architectural framework for a Semantic Grid?

To keep our feet on the ground, the project is developing an architectural framework based on the emerging Open Grid Service Architecture (OGSA)¹⁴ and designed against two case studies from our applications in international insurance settlement and satellite data management: a

⁸ L. Chen, N.R. Shadbolt, C.A. Goble, F. Tao, S.J. Cox, C. Puleston, P.R. "Towards a Knowledge-based Approach to Semantic Service Composition" *2nd International Semantic Web Conference*, 20-24 October, 2003, Sanibel Island, Florida, USA.

⁹ J. D. Myers, C. Pancerella, C. Lansing, K. L. Schuchardt, B. D. "Multi-Scale Science: Supporting Emerging Practice with Semantically Derived Provenance," *Proceedings of the Workshop on Semantic Web Technologies for Searching and Retrieving Scientific Data* at Sanibel Island, Florida on October 20, 2003.

¹⁰ H. Fu, J. G. Frey. "Semantic description and tracking of analysis of chemical data," *Second International Workshop on the Knowledge Grid and Grid Intelligence*, Beijing, China, 2004, 140-149.

¹¹ S. Bowers and B. Ludäscher "An Ontology-Driven Framework for Data Transformation in Scientific Workflows," *International Workshop on Data Integration in the Life Sciences (DILS'04)*, March 25-26, 2004 Leipzig, Germany, LNCS 2994.

¹² M. Bachler, S. B. Shum, Y. -H. Chen-Burger, J. Dalton, D. De Roure, M. Eisenstadt, J. Frey,, J. Komzak, D. Michaelides, K. Page, S. Potter, N. Shadbolt, A. Tate, "Collaboration in the Semantic Grid: a Basis for e-Learning," *Grid Learning Services (GLS2004) at the 7th International Conference on Intelligent Tutoring Systems Workshop (ITS'2004)*, (Maceio, Brazil, 2004), 1-12.

¹³ <http://www.ontogrid.net/>

¹⁴ Foster, H. Kishimoto, A. Savva, D. Berry, A. Djaoui, A. Grimshaw, B. Horn, F. Maciel, F. Siebenlist, R. Subramaniam, J. Treadwell, J. Von Reich, "The Open Grid Services Architecture," <http://www.ggf.org/documents/GFD.30>, 2005.

Virtual Organisation Management System (VOMS) and intelligent debugging. Our first experiment is on a Semantically Aware VOMS, due in mid 2006. Most of the work has focused on a Reference Architecture for the Semantic Grid.

A principled approach to Semantic-OGSA

Currently the Semantic Grid lacks a Reference Architecture or any kind of systematic framework for designing Semantic Grid components or applications. OGSA aims to define a core set of capabilities and behaviours for Grid systems.¹⁴ OntoGrid extends OGSA by explicitly defining a lightweight mechanism that will allow for the explicit use of semantics and defining the associated knowledge services to support a spectrum of service capabilities. Semantic-OGSA (S-OGSA) is guided by seven design principles identified by the project (Figure 3):

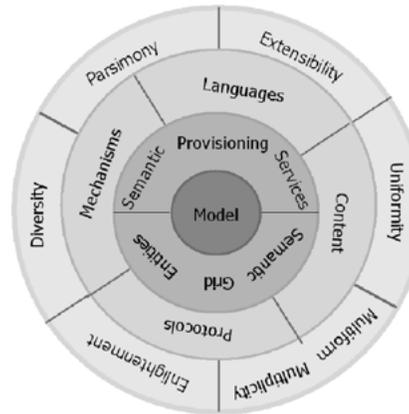


Figure 3: Semantic-OGSA. The outer ring gives the principles; the next ring the main architecture aspects; the next the two major classes of OntoGrid service capability and at the centre is the S-OGSA model.

1. **Parsimony:** the architectural framework should be as lightweight as necessary, minimise the impact on legacy Grid infrastructure and tooling, and not dictate the definition of the contents of the descriptions – these will be application or middleware dependent.
2. **Extensibility:** rather than define a complete and generic architecture, define an extensible and customisable one. Generality is the enemy of applicability.
3. **Uniformity:** Semantic Grids are Grids, so all knowledge services are OGSA-compliant Grid services, and semantic descriptions have a lifetime and a life cycle just like other Grid entities. As metadata stores and ontology services are just special kinds of data services, we have adopted the OGSA-Data Access and Integration specification¹⁵ for their deployment and can potentially exploit other data grid capabilities.
4. **Diversity:** a dynamic ecosystem of Grid services ranging over a spectrum of semantic capabilities will coexist at any one time. Semantic capability may be possible for some Grid resources all of the time, all Grid resources some of the time, or not all resources all of the time.
5. **Multiform + Multiplicity:** the same semantic description may be captured in many representational forms (text, logic, ontology, rule) and any resource's property may have many different descriptions.
6. **Enlightenment:** services should have a straightforward migration path that enables them to become knowledgeable and minimise the cost of doing so.
7. **Conceptual:** S-OGSA is a reference architecture. Thus it should apply equally to different Grid middleware platforms such as the Globus Toolkit,¹⁶ the EU EGEE gLite platform,¹⁷ the UK Open Middleware Infrastructure Institute Release,¹⁸ or regular Web Services.

¹⁵ OGSA Data Access and Integration. Middleware to assist with access and integration of data from separate data sources via the grid.

¹⁶ <http://www.globus.org/>

¹⁷ <http://public.eu-egee.org/>

¹⁸ <http://www.omii.ac.uk/>

These principles pervade OntoGrid development and our thinking.

OntoGrid: A Semantic Grid Reference Architecture

Models, Capabilities and Mechanisms

S-OGSA has three main aspects: the *model* (the elements that it is composed of and its interrelationships), the *capabilities* (the services needed to deal with such components) and the *mechanisms* (the elements that will enable communication when deploying the architecture in an application).

S-OGSA Model. Although there is no standardized overall model of the Grid and its basic concepts, there is a vocabulary associated with OGSA, and there are project specific models^{3,19} and capability focused models like the Common Information Model (CIM)²⁰ from the Distributed Management Task Force and the Job Submission Description Language²¹ from Global Grid Forum. S-OGSA introduces the notion of Semantics into the model of the Grid defining *Grid Entities*, *Knowledge Entities* (e.g. ontologies, rules, text), *Semantic Bindings* between these two for a Grid Entity to become *Semantic Grid Entities*. Semantic Bindings are (possibly temporary) metadata assertions on Grid entities and are Grid resources with their own identity, manageability features and metadata.

S-OGSA Capabilities. S-OGSA is a mixed economy of these semantically enabled and disabled services. We add to the set of capabilities that Grid middleware should provide to include the Semantic Provisioning Services and Semantically Aware Grid Services (Figure 4).

Semantic Provisioning Services dynamically provision an application with semantic grid entities in the same way a data grid provisions an application with data. The services support the creation, storage, update, removal and access of different forms of Knowledge Entities and Semantic Bindings. *Ontology services* store and provide access to the conceptual models representing knowledge; *reasoning services* support computational reasoning with those conceptual models; *metadata services* store and provide access to semantic bindings and the *annotation services* generate metadata from different types of information sources, like databases, services and provenance data. These four build on the past work of members of the consortium: a knowledge parser for extracting information from online sources;²² a metadata store;²³ and a suite of ontologies and supporting tools to generate semantic descriptions for Grid Services.²⁴

Semantically Aware Grid Services exploit knowledge technologies to deliver their functionality, for example metadata aware authentication of a *given identity* by a Virtual Organisation manager service or execution of a *search request* over entries in a semantically enhanced resource catalogue. Sharing this knowledge brings flexibility to components and increases interoperability. OntoGrid is working on a principled re-factoring strategy for legacy Grid Services to quantify the impact on current Grids.

S-OGSA Mechanisms. The model and capabilities are platform independent. To demonstrate the approach in practice, we map the conceptual design to a specific software platform, namely the Globus Toolkit 4, by mapping the semantic bindings to resource properties defined using the Web Service Resource Framework and incorporating S-OGSA entities into the Resource Model of the Common Information Model.

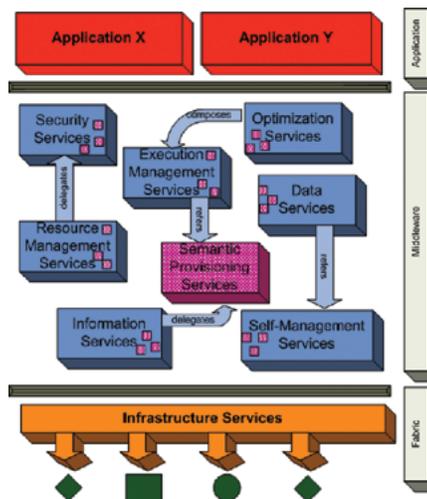


Figure 4: The S-OGSA semantic provisioning services positioned in the OGSA services

¹⁹ N. Sharman, N. Alpdemir, J. Ferris, M. Greenwood, P. Li and C. Wroe, "The myGrid Information Model," *Proceedings of UK e-science All Hands Meeting*, 2004, available from <http://www.mygrid.org.uk/>

²⁰ Common Information Model (CIM) A common definition of management information for systems, networks, applications and services. <http://www.dmtf.org/standards/cim/>

²¹ Job Submission Description Language <http://forge.gridforum.org/projects/jsdl-wg/>

²² Knowledge Parser <http://www.isoco.com/en/innovation/applications/kp.html>

²³ Z. Kaoudi, I. Miliaraki, S. Skiadopoulos, M. Magiridou, E. Liarou, S. Idreos, and M. Koubarakis, "Specification and Design of Ontology Services and Semantic Grid Services on top of Self-organized P2P Networks." *OntoGrid Deliverable D4.1*, 2005.

²⁴ C. Goble, A. Gómez-Pérez, R. González-Cabero, M. S. Pérez. "ODESGS Framework, Knowledge-based markup for Semantic Grid Services," *Proceedings of the Third International Conference on Knowledge Capture (K-CAP 2005)*, Banff, Canada, 2005, 199:200.

Semantic Grid Challenges

Grid Services currently deal with this semantic infrastructure in ad-hoc and hidden ways, providing poor mechanisms for sharing and openly processing knowledge. This makes the knowledge hard to share, and hard to 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, customisation and adaptation difficult, and dependent on scarce human effort. The Semantic Grid aims to provision a semantic infrastructure for Grid infrastructure to improve sharing, enable unanticipated reuse of resources, support interoperability and enable more flexible and configurable middleware.

OntoGrid is a step towards the Semantic Grid. There are many challenges to explore. Many are technical—architectural or theoretical foundations, the maturity of semantic and grid technologies, their appropriateness for the required tasks, their scalability, the separation of grid level and application specific semantics, and making it easier not harder by combining semantic infrastructure with Grid computing infrastructure. Others are operational—gathering and maintaining the semantic content, reliance on unavailable tooling, and convincingly showing the added value of semantics when the return on investment may come downstream, be long term and benefit developers other than the originators. Some are sociological and political—the interplay between the Semantic and the Grid communities, the inter-factional battles within those communities and the legal, security and privacy implications of clearly exposed metadata and automated reasoning.

Acknowledgements The OntoGrid Consortium: Universidad Politécnica de Madrid, Spain (Co-ordinator), The University of Manchester, UK, The University of Liverpool UK, Technical University of Crete (TUC), Greece, Intelligent Software Components, Spain, Y'all B.V., The Netherlands, Deimos Space, S.L, Spain, Boyd International, B.V. The Netherlands. This work is supported by the EU FP6 OntoGrid project (STREP 511513) funded by the Grid-based Systems for solving complex problems.

Glossary

BIRN	Biomedical Informatics Research Network. An NIH initiative supporting distributed collaborations in biomedical science. http://www.nbirn.net/
CIM	Common Information Model. A common definition of management information for systems, networks, applications and services. http://www.dmtf.org/standards/cim/
CMCS	Collaboratory for Multi-scale Chemical Science. Project supporting collaboration through the use of adaptive infrastructure. http://cmcs.ca.sandia.gov/
DFDL	Data Format Description Language. A language for describing the structure of binary and character encoded (ASCII/Unicode) files and data streams. http://forge.gridforum.org/projects/dfdl-wg/
EGEE	Enabling Grids for E-SciencE. EU funded project building grid infrastructure for scientists. http://public.eu-egee.org/
GGF	Global Grid Forum. The community of users, developers, and vendors leading the global standardization effort for grid computing. http://www.ggf.org/
gLite	A lightweight middleware framework from the EGEE project. http://glite.web.cern.ch/glite/
GT(4)	Globus Toolkit (4). An open source software toolkit used for building Grid systems and applications. Developed by the Globus Alliance. http://www.globus.org/toolkit/
JSDL	Job Submission Description Language. https://forge.gridforum.org/projects/jsdl-wg/
Matlab	A language and environment supporting computationally intensive tasks. http://www.mathworks.com/
OGSA	Open Grid Services Architecture. A set of core capabilities and behaviours that address key concerns in Grid systems. http://www.globus.org/ogsa/
OGSA-DAI	OGSA Data Access and Integration. Middleware to assist with access and integration of data from separate data sources via the grid. http://www.ogsadai.org/
OMII	Open Middleware Infrastructure Institute. An EPSRC funded initiative providing reliable, interoperable and open-source Grid middleware. http://www.omii.ac.uk/
P2P	Peer to Peer. Architectures which allow autonomous peers to interoperate in a decentralized, distributed manner for fulfilling individual and/or common goals
SAML	Security Assertion Markup Language. A language for exchanging authentication and authorization data between security domains. http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=security
S-OGSA	Semantic OGSA.
VO	Virtual Organisation. Flexible, secure, coordinated resource sharing among dynamic collections of individuals, institutions, and resources.
VOMS	Virtual Organisation Management Service. A service managing a VO.
WSRF	Web Services Resource Framework. A framework defining conventions for modelling and accessing stateful resources using Web services http://www.globus.org/wsrfl/

How to Build an International Grid: Infrastructure, Applications & Community

Introduction

The Enabling Grid for E-sciencE project (EGEE) is Europe's flagship Research Infrastructures Grid project¹ and the world's largest Grid infrastructure of its kind. It involves more than 70 partners from 27 countries, arranged in 12 regional federations, and providing more than 16,000 CPUs, more than 160 sites and 10 petabytes of available network storage. This infrastructure supports six scientific domains and more than 20 individual applications.

Started in March 2004, EGEE has rapidly grown from a European to a global endeavour, and along the way learned a great deal about the business of building production-quality infrastructure. The consortium behind this effort represents a significant proportion of Europe's Grid experts, including not only academic intuitions but also partners from the Research Network community and European industry. This article outlines the project's structure and goals, its achievements and the importance of cooperation in such large scale international efforts.

A distributed effort – project structure and goals

The aim of EGEE is to leverage the pre-existing grid efforts in Europe, thematic, national and regional, to build a production quality multi-science computing Grid. As a result, the primary objective is to build the infrastructure itself, connecting computing centres across Europe (and more recently, around the globe) into a coordinated service capable of supporting 24/7 use by large scientific communities. To support this production service, the project also aims to re-engineer existing middleware components to produce a service-oriented middleware solution. Finally, the project aims to engage the maximum number of users running applications on the infrastructure through dissemination, training and user support. These tasks have been divided into different activity areas, which are tackled by different groups within the project. These groups are distributed across a number of partner institutes with relevant experience, such that the project helps to connect its partners and encourage knowledge transfer in the process of achieving its goals.



Fig 1: Extent of EGEE infrastructure for EGEE-II

Building a large, secure, stable and scalable infrastructure is perhaps the key feature of EGEE. From the start, the project benefited from the resources of the international High Energy Physics (HEP) community, leveraging these to build a Grid infrastructure for all scientific disciplines. These HEP resources come from the computing systems built to support the forthcoming Large Hadron Collider (LHC) being built at CERN² in Switzerland. More specifically, EGEE formed a strategic alliance with the LCG³ (LHC Computing Grid) project, which independently was deploying an international distributed computing infrastructure.

Connecting and sharing – growing a global infrastructure

Building a large, secure, stable and scalable infrastructure is perhaps the key feature of EGEE. From the start, the project benefited from the resources of the international High Energy Physics (HEP) community, leveraging these to build a Grid infrastructure for all scientific disciplines. These HEP resources come from the computing systems built to support the forthcoming Large Hadron Collider (LHC) being built at CERN² in Switzerland. More specifically, EGEE formed a strategic alliance with the LCG³ (LHC Computing Grid) project, which independently was deploying an international distributed computing infrastructure.

Fabrizio Gagliardi

EGEE Project Director – CERN

Bob Jones

EGEE Technical Director – CERN

Owen Appleton

EGEE Communications Officer – CERN

¹ EGEE is funded by the EU Information Society & Media directorate through the Sixth Framework Programme, contract number INFSO-RI-508833.

² The European Nuclear Research Organization, <http://www.cern.ch/>

³ <http://lcg.web.cern.ch/LCG/>

With an infrastructure of a considerable size from the HEP community available from day one, EGEE has been able to concentrate on delivering a working infrastructure, with a main production service supported by pre-production, testing and development services, and even specialised infrastructure for dissemination and training.⁴ This initial pool of resources supplied by the HEP community helped to encourage the other pilot application domains, the Biomedical science community, to contribute their own resources and run their own production challenges, thus encouraging other domains to join the project.

It also became clear during the early part of the project that restricting this effort to Europe made little sense given the distributed nature of many scientific communities and the large number of resources, both in terms of knowledge and hardware, in other parts of the globe. EGEE began to extend its efforts beyond its original partners early on in the project through extension of the infrastructure into South-Eastern Europe through the SEEGRID⁵ project and into digital library applications through the DILIGENT⁶ project. This successful policy of collaboration and extension has continued, with EGEE building relationships with major sister projects in areas such as the United States (OSG) and Asia (NAREGI), as well as through support for related projects (such as BalticGrid, EUChinaGrid, EELA and EUMedGrid) that extend the EGEE infrastructure to new geographical areas. Such associations are an important part of EGEE's role as an incubator, both within Europe and beyond, actively supporting a wide range of Grid efforts, from infrastructure to application projects. Through these projects, EGEE has spread the knowledge it has accumulated in all areas of its work, from making applications Grid compliant to managing infrastructure. This cooperative spirit is also represented in the way that the infrastructure is managed. Initially run from a central centre at CERN to spread both the workload and the knowledge generated by managing large scale infrastructures, responsibility now rotates around centres across Europe (with future plans for centres in the US and Asia)

⁴ EGEE uses GILDA, a dedicated testbed for dissemination and training provided by Italy's Istituto Nazionale di Fisica Nucleare. <https://gilda.ct.infn.it/>

⁵ South Eastern European Grid-enabled eInfrastructure Development, <http://www.see-grid.org/>

⁶ A Digital Library Infrastructure on Grid Enabled Technology, <http://www.diligentproject.org/>

Re-engineering and integration – producing modern middleware

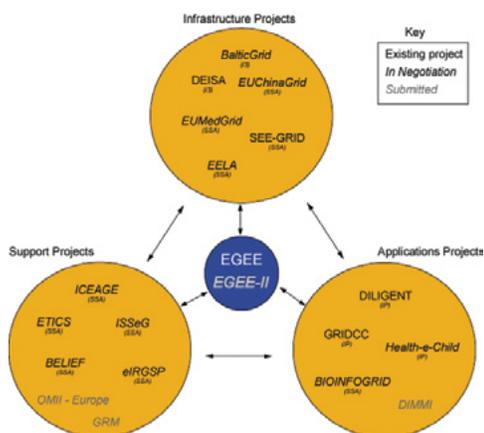


Fig 2: Projects related to EGEE and EGEE-II

The resulting software stack provides a full set of Grid foundation services, as well as a range of higher level services.

This modular approach, combines best-of-breed elements from other middleware sources, allowing outside projects interested in the middleware to install only the elements of gLite relevant to them, developing their own specialist, high-level services on top of it. The gLite stack also contains the foundations necessary for interoperability with other Grid systems, in

From its inception, EGEE had considerable advantages in the area of middleware. EGEE is in many ways the successor to the European DataGrid project,⁷ which had previously developed a well built middleware solution, the EDG middleware. This stack had already been further developed by the LCG project into the LCG-2 middleware, providing EGEE with a working middleware stack to deploy from day one. In parallel, EGEE has also developed a new middleware solution, gLite,⁸ tailored to the multi-science user communities it supports. Rather than starting from scratch, gLite takes components from a large number of software sources, re-engineering some of them and integrating them into a modern, lightweight, service-oriented middleware solution.

⁷ <http://eu-datagrid.web.cern.ch/eu-datagrid/>

⁸ Pronounced "gee-lite", <http://www.glite.org/>

How to Build an International Grid: Infrastructure, Applications & Community

particular in the area of security frameworks, and its development team actively participates in security working groups through the Global Grid Forum⁹ and other such bodies.

⁹ <http://www.gridforum.org/>

The gLite stack is released under a permissive, business-friendly open source license, which facilitates and encourages its use by outside groups. With gLite in use within the project, outside groups such as the DILIGENT digital library project are already making use of gLite on their own Grid infrastructure, and it is hoped that more groups, and eventually industry, will join them in the future.

Prototype to applications – infrastructure in action

EGEE began with two pilot application domains, High Energy Physics (HEP) and Biomedical Science. The HEP domain includes close collaboration with LCG to process data from the international LHC experiment communities, but also includes applications from other HEP projects such as CDF, D0, Zeus and Babar. In the Biomedical Science field, some 10 different applications are already running, ranging from protein sequence analysis to molecular docking studies used to look for new treatments for Malaria.

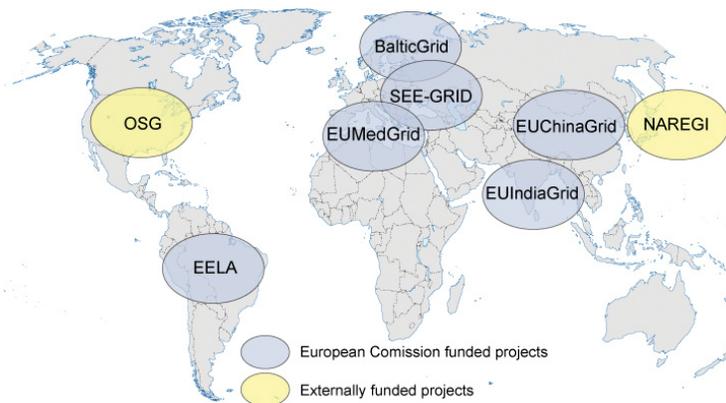


Fig 3: EGEE-II geographical extension through related projects

In addition to these pilot domains, a number of other groups have joined the EGEE infrastructure since the project started, namely Computational Chemistry, Astrophysics, Earth Sciences and Geophysics. Such new groups can join the infrastructure through a system called the “Virtuous Cycle.” In this process, new communities become aware of the availability of the EGEE service through outreach events, personal contacts or through contacts with project members in their local area and can try the grid through online demonstrations. Following this, they interact with local resource centres, which provide access to resources and aid in porting applications to the EGEE infrastructure. This allows new applications to come from within the project in an organic manner and be identified with a nearby group able to communicate the new application’s requirements to the rest of the project. Once on the Grid, new application groups receive training in all the appropriate skills they need in order to make them a self supporting community. Finally, the new group becomes an established user community on the EGEE infrastructure, demonstrating to other potential users the benefits of Grid technology and encouraging them in turn to get involved with EGEE.

The vibrant and extensive user community formed from users in these application domains is in many ways EGEE’s greatest achievement. No other production grid infrastructure exists of this size or with this breadth of active users. Growing since the start of the project, the number of successful jobs per day on the infrastructure had exceeded 19,000 by June 2005. EGEE is not

only breaking new ground in understanding the unique challenges that running such a truly interdisciplinary infrastructure presents, but also passing this knowledge on to sister projects in other parts of the world, industry and more focussed Grid projects.

Apart from the various academic scientific communities that are involved with EGEE, the project also supports an industrial application from French firm Compagnie General de Geophysique (CGG), who support the EGEODE Virtual Organisation,¹⁰ used for basic geophysics research. EGEODE benefits EGEE, CGG and the geophysics community in general by freely distributing the results of its research, as well as helping EGEE attune itself to industrial requirements and expectation for the future of Grid computing as a commercial service.

¹⁰ Virtual Organisations (VOs) are systems for allowing distributed communities to work together and share resources on a Grid infrastructure.

Toward a permanent research infrastructure

EGEE was originally conceived as the first two years of a four year programme and, in keeping with this vision, the consortium behind EGEE recently submitted a proposal to the recent EU Information Society Research Infrastructures funding call for the second half of this programme, the EGEE-II project.

EGEE-II is a further elaboration of the EGEE mission, learning from the experience of the previous project and featuring a considerably expanded consortium and refocused mission. As well as increasing its consortium to over 90 partners from 32 countries, it increases its global vision by formalising relationships with partners in the USA, Taipei and Korea. In the USA, this also includes other large scale Grid projects such as OSG and Grid3, allowing both sides to profit from one another's experience. Further extension of the infrastructure to the Baltic, Mediterranean area, China and Latin America will be achieved through related projects also submitted to EU Information Society funding calls.

Since the beginning of EGEE, Grid technology has matured considerably, with a great number of projects across the globe producing interesting results. EGEE-II has been planned in light of these developments, allowing it to profit from them as well as passing information and experience back into the community. This has led to a refocusing of the project activities, with a greater emphasis on infrastructure management and a new dedicated effort in middleware certification integration and testing. In parallel, middleware re-engineering within the project will focus more on integrating components from outside sources including Globus, Condor, and the Virtual Data Toolkit (VDT) and from related European Grid projects.

In the applications area, EGEE-II will continue to increase the number of scientific domains and applications running on the infrastructure. This will notably include collaboration with the International Thermonuclear Experimental Reactor project (ITER)¹¹ on fusion applications, as well as support for any other interested partners.

¹¹ <http://www.iter.org/>

Overall, the long term goal of EGEE and EGEE-II is to establish a permanent public Grid infrastructure to support research of all types. Through the course of EGEE, it has become clear that profiting from such infrastructures requires the greatest possible level of interconnection with other similar efforts. As a result, through the course of EGEE, such collaboration has increased considerably, and the plans for EGEE-II were framed with such collaboration in mind. Through this strategy, not only is the effectiveness of the individual projects and infrastructures improved, but it promotes common standards and interoperability crucial to the future of Grid technology for both academic and industrial users.

SIMDAT

Introduction

In the context of this project, a Grid is defined to be a software system that provides uniform and location independent access to geographically and organizationally dispersed, heterogeneous resources that are persistent and supported. Typically, these shared assets are under different ownership or control. The SIMDAT project¹ is developing generic Grid technology for the solution of complex application problems and demonstrating this technology in several representative industry sectors. Special attention is being paid to security, e.g. where third-party suppliers have need-to-know access to data, and correlation and inference may provide insight into confidential processes. The objective is to accelerate the uptake of Grid technologies in industry and services, provide standardised solutions for capability currently missing, and validate the effectiveness of a Grid in simplifying processes used for the solution of complex, data-centric problems.

The SIMDAT consortium is comprised of leading software and process system developers—IBM, IDESTYLE Technologies, InforSense, Intel, Lion Bioscience, LMS International, MSC Software, NEC, Ontoprise and Oracle; Grid technology specialists—Fraunhofer Institute AIS, Fraunhofer Institute SCAI, IT Innovation, Universitat Karlsruhe, Universite libre de Bruxelles and the University of Southampton; and representatives from strategic industry and service sectors—Audi, BAESystems, DWD, EADS, ESI, EUMETSAT, ECMWF, GlaxoSmithKline, MeteoFrance, Renault, and the UK Met Office. IT Innovation is leading the basic Grid infrastructure level architecture work in SIMDAT and this article will therefore be focused on this aspect of the project rather than the applications.

Grids for complex problem solving in industry

Development of industrial and large-scale products and services poses complex problems. The processes used to develop these products and services typically involve a large number of independent organisational entities at different locations grouped in partnerships and supply chains. Grid is connectivity plus interoperability and is a major contributor to improved collaboration and an enabler of virtual organisations. It has the potential to substantially reduce the complexity of the development process, thereby improving the ability to deal with product complexity.

The heart of the issue is data. Applications and their associated computing power are central to the product development process. Grid technology is needed to connect diverse data sources, to enable flexible, secure and sophisticated levels of collaboration and to make possible the use of powerful knowledge discovery techniques.

Key to seamless data access is the federation of problem-solving environments using grid technology. The federated problem solving-environments will be the major result of SIMDAT. Seven key technology layers have been identified as important to achieving the SIMDAT objectives:

- an integrated grid infrastructure, offering basic services to applications and higher-level layers

Mike Boniface and
Colin Upstill

University of Southampton IT Innovation
Centre

¹ <http://www.simdat.org/>

- transparent access to data repositories on remote Grid sites
- management of Virtual Organizations
- workflow
- ontologies
- integration of analysis services
- knowledge services.

The strategic objectives of SIMDAT are to:

- test and enhance data grid technology for product development and production process design
- develop federated versions of problem-solving environments by leveraging enhanced grid services
- exploit data grids as a basis for distributed knowledge discovery
- promote de facto standards for these enhanced grid technologies across a range of disciplines and sectors
- raise awareness of the advantages of data grids in important industry sectors.

SIMDAT focuses on four exemplar application areas: product design in the automotive, aerospace and pharma industries; and service provision in meteorology. For each of these application areas a challenging problem has been identified that will be solved using Grid technology, e.g. distributed knowledge discovery to enable better understanding of the different Noise, Vibration and Harshness (NVH) behaviour of different designs of cars based on the same platform; Grid technology will allow seamless access to all relevant data for all engineers of the development centers of large multinational car manufacturers.

SIMDAT Architecture

All application sectors deploy existing problem-solving environments for product and process design. Each application activity in SIMDAT is integrating Grid middleware into existing applications to provide a demonstration of distributed, collaborative work in complex problem solving. The vendors of these environments require an acceptable level of stability in middleware technology before it will be adopted with their products and delivered to customers. Well-designed and accepted standards are essential for technology uptake.

Examining Grid infrastructure state-of-the-art, it is clear that even the core technology, which underpins higher-level services such as resource and execution management, is still evolving. In the future, core features should be part of a standards-compliant architecture, so application developers can use them more easily and so they can choose between different interoperable Grid implementations.

The Open Grid Services Architecture² (OGSA) represents an evolution towards a Grid system architecture based on Web services concepts and technologies. OGSA Profiles for various higher-level functions are beginning to be developed but there are certainly no OGSA compliant grid implementations. Even the underlying proposed standard WS-Resource Framework³ (WS-RF) is still somewhat controversial and has yet to prove its value. Therefore, the challenge of standardising the Grid programming model and associated management services is still unfulfilled.

² <http://www.globus.org/ogsa/>

³ <http://www.globus.org/wsrf/>

SIMDAT

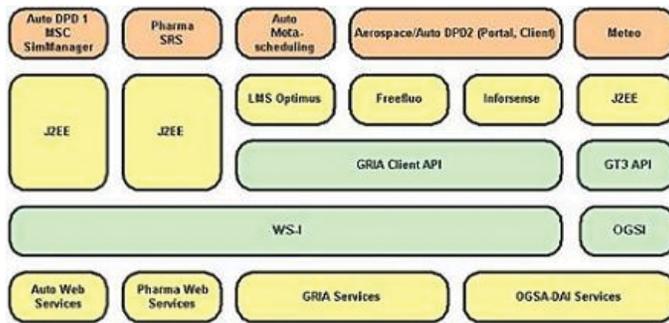


Figure 1: Initial SIMDAT Architecture

SIMDAT has adopted a pragmatic approach for using existing Grid infrastructure and Web Service technologies. The application sectors faced the challenge of selecting Grid technologies that best fitted their scenarios, even if they did not provide all of the necessary functionality. Key among these are GRIA⁴ (open source Grid middleware developed by IT Innovation to enable commercial use of the Grid in a secure, interoperable and flexible manner), Globus,⁵ and J2EE⁶ portals. The application activities concluded that in the short-term, unless and until a standardised Grid programming model is agreed upon, new developments should be based on Web Service standards such as WS-I.⁷ GRIA emerged as a core technology to support collaborative work in the SIMDAT aerospace and automotive activities because of its availability, its adherence to WS-I, and its explicit support for B2B collaborations. The SIMDAT meteorology activity is developing a Data Grid based on Open Grid Services Architecture Data Access and Integration⁸ (OGSA-DAI) and Web Service technology, while the SIMDAT pharmaceutical activity is deploying a Web Services Grid leveraging E2E (end-to-end) security component developed during the GEMSS project.⁹ The initial architecture (Figure 1) shows how WS-I can provide a common API for distributed services, but it does not currently meet Grid infrastructure requirements for providing a standardised approach for managing stateful resources, as proposed by WS-RF.

SIMDAT provides application sectors with a Grid infrastructure roadmap that tracks the rapidly changing Grid landscape. During 2005, the situation has evolved significantly as GRIA continues to be developed and as technologies such as GT4¹⁰ and gLite¹¹ emerge. Following the recent delivery of the first SIMDAT prototypes, the application activities are feeding back lessons learned, which will be factored into the roadmap and into the Grid technologies. In the longer term SIMDAT aims to achieve interoperability between different Grid infrastructures such as GRIA and GT4, with a Grid service API based on WS-RF, although the level of compliance may differ between implementations.

Aerospace Case Study

The aerospace industry deals with highly complex products that have data creation, management and curation requirements that span hundreds of collaborating organisations over a 50-year lifecycle. Partners in a product team need to collectively manage thousands of inter-related processes and this leads the industry to expend considerable time and effort in the access, transmission, control, translation and sharing of data.

The primary focus of the aerospace activity in SIMDAT is the development and deployment of existing and emerging Grid technologies and concepts to enhance the collaborative engineering

⁴ <http://www.gria.org/>

⁵ <http://www.globus.org/>

⁶ <http://java.sun.com/j2ee/index.jsp/>

⁷ <http://www.ws-i.org/>

⁸ <http://www.ogsadai.org.uk/>

⁹ <http://www.gemss.de/>

¹⁰ <http://www.globus.org/toolkit/>

¹¹ <http://glite.web.cern.ch/glite/>

of sophisticated products. The improvement in the ability to handle complex problems is not delivered simply through the connectivity that Grid offers, but through the deployment of industry-strength middleware and advanced ontology-based techniques to radically improve the efficiency of the data exchange, both between applications and between organisations

The initial aerospace deployment simulates the multi-disciplinary collaborative configuration design of a low-noise, high-lift landing system. The scenario is typical of sub-system design problems in the context of future-concept, unmanned cargo vehicles that need to use airfields in noise-sensitive locations. The scenario has been designed to show how Grid technologies can support the aggregation of distributed capabilities operating across organisational boundaries.

The deployment of the aerospace prototype demonstrates how Grid technologies can support pan-European inter-Enterprise collaborative development of complex products (Figure 2). Each organisation within the aerospace deployment operates as a GRIA service provider offering specialised engineering services such as optimisation (University of Southampton), parameterised CAD generation (University of Southampton), aerodynamics (BAE SYSTEMS) and aero-acoustics (EADS). GRIA's explicit business process support for dynamic, bi-lateral QoS agreements allows project managers at aerospace companies to create inter-Enterprise multidisciplinary design teams in a secure, managed, auditable and accountable environment.

GRIA has been significantly enhanced to support the aerospace application scenario through integration with other key Grid technologies. OGSA-DAI WS-I has been integrated with GRIA to provide distributed data access for relational data and simulation files. The GRIA OGSA-DAI service provides security and enforces a business model for managing distributed data resources. A GRIA workflow service has been developed, based on IT Innovation's open source workflow enactor, Freefluo.¹² This allows aerospace engineers to publish workflows as services that can be executed by distributed clients.

The Grid programming environment is provided by the Taverna¹³ workbench, which has been enhanced to support GRIA's business process and WS-I Basic Security Profile. Aerospace engineers integrate legacy applications as Grid services using GRIA and compose these applications into workflows using Taverna. The workflows can then be published to GRIA's workflow enactment service, allowing clients to compose hierarchical distributed workflows that cross organisation boundaries. At the lowest level, the aerospace workflows consist of a simple computational sequence of meshing, solving and post-processing. At higher-levels, the workflows are much more complex. For example, the design workflow (Figure 3) explores the design space by creating a design of experiments (DoE) from a given input specification and iteratively calculating the results for each design point in the DoE. The workflow implementation invokes optimisation, CAD generation and compute workflows, along with staging input and output data in a design database accessed through OGSA-DAI.



Figure 2: Pan-European deployment of Grid technologies in Aerospace

¹² <http://freefluo.sourceforge.net/>

¹³ <http://taverna.sourceforge.net/>

The Next Generation Grid

Do we really need a next generation of the Grid?

Mark Parsons
NextGRID Project Chairman

To some people it seems premature to talk of the next generation of the Grid when in many cases the Grid has yet to deliver according to its original vision. Grid research has come a long way since it was originally mooted—in terms analogous to the electric power grid—as an infrastructure that was always-on and delivered chargeable access to compute, data and other resources when and wherever they were required. Pioneering projects, largely science-based, in Europe, the US and Asia have demonstrated the positive benefits afforded by large-scale, widely distributed computation and data access and such projects are now undertaking previously impracticable scientific research. This is particularly true in the health sector where some large cancer research projects are now gathering speed and will hopefully afford real benefits and breakthroughs across society.

However, although the Grid can be said to be delivering in a scientific context, the same is not true in the business domain. Visit any investment bank in the United Kingdom and they will (privately) talk proudly of the success of *their* Grid. In reality, they are actually talking about the success of their clustered computing approach. There are two main reasons for this. Firstly, the hijacking of the “Grid” word by over-eager vendor marketing departments following the dot com bubble in the early part of this decade has confused many potential users about what the Grid is really for—inter-enterprise, joined-up computing. Secondly, and more importantly, the Grid used and promulgated by the science and research communities does not take into account the typical regulatory and management issues faced by many industries. Unless the Grid can be seen to offer real benefits to business it will remain a powerful tool for science and will be largely ignored by business, except in its simplest application server and clustered computing form. In the worst case we will see a complete divergence in Grid computing between science and business.

It doesn't have to be like this.

It is very easy to complain that the Grid to date has failed to link its developments to the real needs of its users. In the scientific domain this simply is not true. Wide-ranging requirements-gathering activities have taken place and will continue. These activities have helped to guide the development of the tools most needed by these programmes of scientific research. In most cases these are programmes of research where a specific end-point is reasonably clear, and the main motivation for using Grids is to collaborate in order to pool resources. In the business domain, the requirements that Grids have to meet are far broader and more varied. A wide variety of projects, notably in the UK and Europe, have been undertaken, and there have been many notable demonstrations of the efficacy of Grids both in the cluster and broader Grid contexts. However, these projects did not produce universal solutions spanning many business applications, because different solutions were required in each case. For example, the GRASP project used “traditional” academic Grid principles to support resource sharing within a cooperative of application service providers, providing higher performance and reliability for ASP services, but requiring mutual trust between the providers. The GRIA project implemented an inter-enterprise collaboration infrastructure allowing the users to pool resources obtained on commercial terms from independent service providers. The GEMSS project took a similar approach for medical simulation services, but resources from different service providers cannot

The Next Generation Grid

be pooled in a single application because that would make it very difficult to meet European privacy regulations for processing patient data.

All of these commercial Grid prototypes support the application requirements typically found in academic Grids, but they all had to be specialised in some way to meet specific business requirements in the sector or scenarios they were designed to support. There remains a lack of consensus on what constitutes a usable business Grid, and without this, the impact of Grids on business will remain very limited. To overcome these barriers, it is not enough to simply implement solutions that meet the business requirements of individual users. The World Wide Web was not delivered by studying the requirements of business – but it continues to be truly transformational in the way all of us live and work. The next generation of Grid must therefore be truly transformational. It must go well beyond the stated requirements of science and business and it must also be prepared to challenge current orthodoxies. Those of us who believe in the Grid understand that, as it delivers over the next decade, we will begin to see the true value of this computing revolution.

Getting from where we are now to where we want to be will not be easy. Neither is it easy to visualise that end point clearly. We just know it will be there. In some business sectors, enormous resistance to change has built up and the Grid is perceived to have failed to deliver or it has delivered in a very constrained way – such as cluster computing. For instance, a common refrain in the financial services sector is *“we’ll never deploy wide-area Grids – they’re alright for scientists but what about our regulatory and security requirements?”* The purpose of the NextGRID Project is to challenge some of these attitudes and to undertake the necessary thinking to make the Grid truly able to deliver.

NextGRID,¹ funded by the European Union’s Information Society & Media directorate, is a three year, €16.5million research project with partners drawn from across European academia and business. With 13 of the 22 partners coming from the business domain, including SAP, Microsoft, Fujitsu, BT and NEC, the project has a strong focus on tackling the “Grid for business” agenda. The project is currently the only project worldwide that is specifically focused on driving the architecture of the Grid forward.

¹ <http://www.nextgrid.org/>

NextGRID challenges

Information and communication technologies are recognised as having a key role in Europe’s transformation into a dynamic, competitive knowledge-based economy. Sustained success is increasingly reliant on flexibility in business processes, which allows businesses to adapt to a changing global environment. IT applications and services are an essential enabler for this flexibility. Largely to meet this need there is an ongoing clear shift in the market towards a service-oriented approach to IT systems. This is allowing consumers to obtain a wide range of services as required from a range of providers, delivered via a ubiquitous telecommunications infrastructure. The emergence of this infrastructure has allowed users to enjoy permanent global connectivity from a range of wired and wireless devices without needing to be concerned with the technologies and networks involved.

The Grid has the potential to make a significant advance beyond the World Wide Web, by turning it from a passive information medium into an active tool for creating and exploring new knowledge and thereby fuelling business and industry. Today, as discussed above, this potential is unrealised and, without far more cost-effective and universally applicable technology, will remain so. A crucial missing element is the ability to compose services from independent sources in a standardised and cost-effective way. To go beyond current business use of the Grid, applications should be capable of executing on an inter-enterprise Grid infrastructure.

Current Grid systems do not address this service composition challenge—they impose business models on users and application developers, usually based on the “traditional” virtual organisation model for collaboration between mutually trusting parties. Until the Grid can support a wide range of dynamically evolving business models, while maintaining stability as seen by each stakeholder, it is hard to see how the Grid can support third-party application development, which is one of the key drivers behind the success of non-Grid computing platforms.

The separate interests of independent stakeholders cannot be resolved *a priori* as is the case for non-Grid applications designed to execute in a single domain. This implies that a Grid infrastructure must be capable of combining the different business models used by different stakeholders at run time, so the Grid presents a stable interface to each stakeholder. This is of course analogous to the World Wide Web today where a multitude of Web servers and Web browsers (mostly) happily coexist with each other. Furthermore, commercial business models are essential for the Grid’s long-term viability.

NextGRID’s Vision

The broad NextGRID vision is of a networked IT infrastructure to support an unlimited range of applications and business processes throughout their lifecycle. This includes all resources—hardware, software, data and services, available from a complex ecosystem of providers.

The primary goal of NextGRID is to define the architecture that will lead to the emergence of the Next Generation Grid. This will prepare the way for the mainstream use of Grid technologies and their widespread adoption by organisations and individuals from across the business and public domains. In addition to new architectural designs, NextGRID will contribute to the key middleware components, application support mechanisms, know-how and standards that underpin the Next Generation Grid.

Of course, NextGRID cannot address these objectives alone. The participants in NextGRID are the representatives of a much larger community of researchers, technology vendors, service providers and users. We inspire and work with this wider community, providing critical input and thought leadership to the development of the architecture for future Grids, incorporating our results into widely accepted standards, and so encompassing a much larger body of work within our own organisations and in the community at large. We also understand that parts of our work will be incremental and parts revolutionary.

The project structure is built around the architectural design process. This process is informed by the development work, business and operational activities and application experimentation.

The Next Generation Grid

At the end of each six-month design cycle, the results are fed into the development activities, which focus on Grid foundations, dynamics and interactions. The consolidated outputs of the project are exploited up by its standardisation activity and the business partners in the project.

The work of the project is very broad and this article is too short to relate all of it here. The project has now completed two architecture cycles and the broad thrust of our activities has been defined. Rather than detailing our many activities there is more value to be gained by focusing on one particularly important innovation—the idea that Grids can be built up from pairwise inter-enterprise relationships governed by Service Level Agreements (SLAs) that capture the mutual interests of each pair of participants.

NextGRID Architecture and Service Level Agreements

Our initial architectural designs assume that applications will be constructed by composing services, each of which has some common properties and behaviours. When executing applications, we can assume that certain core “infrastructure” services or properties are available in the environment of the application. In the context of NextGRID we are building on the HTTP(S), SOAP, WSDL1.1, WS-Addressing, WS-Security, SAML1.1 and X.509 protocols and the OGSA WSRF Basic Profile 1.0 and SLA Template interfaces. A key architectural requirement is that service composition should result in self-similar structures that are themselves amenable to NextGRID composition rules.

A key aspect of the current NextGRID conceptual architecture is that all interactions will be governed through bipartite “partnership” SLAs. NextGRID believes that SLAs should be used to build relationships between service providers and consumers, and provide the necessary information to set up the environment and components to manage the service. The SLA should outline details that are agreed by both parties, and allow for the service to be operated and monitored in accordance with the consumer requirements and in an economically sustainable manner.

Neither the service provider nor the consumer will gain a significant advantage by violating an SLA. The customer will not get the service they require, and the provider’s reputation will be damaged –perhaps to the stage where the customer will not use the service again. It is therefore proposed to have a framework that is less focused on monitoring of every element of every transaction and see the relationship between provider and consumer as a partnership within a context – with that context being provided by the SLA. Where necessary, monitoring and enforcement are provided for, but the aim of the SLA idea is to focus on the partnership and agreement side rather than the violation side of the contract.

The benefit of such a proposal is that the NextGRID architecture can support both community minded approaches as well as the commercial offering of services. SLAs allow for services to be provided in exchange for an equivalent set of services or a cash purchase. In a commercial context where services are provided for a fee, and the fact the service is provided on a Grid infrastructure is irrelevant to the end user, it is more important to provide specific QoS levels that need to be communicated, agreed upon and upheld.

Approaching the issue of SLAs in this partnership model allows for a lighter-weight monitoring infrastructure and avoids having a monitoring system that is more expensive to provide in economic, computational or time terms than some of the services it is tasked with monitoring.

We believe that an SLA is a key component to be considered at all stages in the lifecycle of a service. The policies for managing the service, the probes for monitoring it and the acceptable quality of service terms to offer to a consumer should be produced at the same time as the service is designed and developed. This ensures the correct information is available to be able to guarantee the QoS levels necessary that a consumer will consider entering into an agreement with a provider to use the service.

Having completed the initial architectural design work with regard to SLAs, the Grid foundations activity is now producing a prototype implementation of an SLA template interface. Other tasks are analysing how “collective” inter-enterprise computing business issues can be addressed through pairwise SLAs between participants, by studying how accountability and billing can be represented in SLA terms, for example. These ideas will then be implemented, so that NextGRID industrial partners can experiment with using this technology for real world applications, to prove its efficacy or, alternatively, show where we have gone wrong in our design and thinking. It is through the many strands of work within NextGRID such as this that we are building the next generation of Grid that will meet the needs of business and commerce.

Conclusion

Gartner Group—the well known IT analysis company—have widely publicised the idea of a “hype cycle for emerging technologies” In the *Gartner Hype Cycle Special Report for 2005*, Grid Computing finds itself halfway down the slope from the “peak of inflated expectations” to the “trough of disillusionment.” Interestingly however, Gartner indicates that they expect it to reach the “plateau of productivity” within 2-5 years.

Achieving our goal of productive Grid computing for business and science requires that we focus both on incremental improvements of current technologies and also that we are prepared to think beyond the status quo and try out new ideas. This is a key premise of the NextGRID project.

Acknowledgments

The NextGRID project is funded by the European Commission under contract number 51563. This paper expresses the opinions of the author and not necessarily those of the European Commission. The European Commission is not liable for any use that may be made of the information contained in this paper.

PERSPECTIVES

The Next Decade in HPC

The Role of High Performance Computing

Craig Mundie
CTO, Microsoft Corporation

I have had a long history in the HPC community: I spent from 1982 to 1992 as the founder and architect of Alliant Computer Systems Corporation in Boston. We spent a long time trying to develop tools and an architecture whose components today would look like they were all fairly slow. But architecturally, many of the concepts that were explored back then by Seymour Cray and the many supercomputer companies—of which Alliant was just one—still to this day represent the basic architectures that are being reproduced and extended as Moore's Law continues to allow these things to be compacted.

In my present role as CTO of Microsoft, it is probably fair to say that I have been the 'god-father' in moving Microsoft to begin to play a role in the area of technical computing. Up to now, the company has really never focused on this area. It is of course true that there are many people in the world who, whether they are in engineering or science, business or academia, use our products like tools on their desktop much like they think of pencil and paper. They would not want to work without them. But such tools are never really considered as an integral part of the mission itself. It is my belief that many of the things that HPC and supercomputing have tended to drive will become important as you look down the road of general computing architectures.

The worldwide aggregate software market in technical computing is not all that large on a financial scale. However, Bill Gates and I, over the last couple years, have agreed that engaging with HPC is not just a question of how big the market is for software per se in technical computing. Rather it is a strategic market in the sense of ultimately making sure that there will be well-trained people who will come out of a university environment and help society solve the difficult problems it will be facing in the future. The global society has an increasing need to solve some very difficult large-scale problems in engineering, science, medicine and in many other fields. Microsoft has a huge research effort that has never been focused on such problems. I believe that it is time that we started to assess some application of our research technology outside of our traditional ways of using it within our own commercial products. We think that by doing so, there is a lot that can be learned about what will be the nature of future computing systems

Many of the things that we thought of as *de rigueur* in terms of architectural issues and design problems in supercomputers in the late eighties and early nineties have now been shrunk down to a chip. Between 2010 and 2020 many of the things that the HPC community is focusing on today will go through a similar shrinking footprint. We will wake up one day and find that the kind of architectures that we assemble today with blades and clusters are now on a chip and being put into everything. In my work on strategy for Microsoft I have to look at the 10 to 20 year horizon rather than a one to three year horizon. The company's entry into high performance computing is based on the belief that over the next 10 years or so, there will be a growing number of people who will want to use these kinds of technologies to solve more and more interesting problems. Another of my motivations is my belief that the problem set, even in that first 10-year period, will expand quite dramatically in terms of the types of problems where people will use these kinds of approaches.

There was a time certainly, when I was in the HPC business, that the people who wrote high performance programs were making them for consumption largely in an engineering environment. Only a few HPC codes were more broadly used in a small number of fields of academic research. Today, it is doubtful whether there is any substantive field of academic research in engineering or science that could really progress without the use of advanced computing technologies. And these technologies are not just the architecture and the megaflops but also the tools and programming environments necessary to address these problems.

Computer Science and the Science and Engineering Community

In parallel with these developments in HPC, we are no longer seeing the kind of heady growth in the number of trained computer scientists produced by the world's universities. In fact, in the United States, this number is actually going down. The numbers are still rising in places like India and China right now, but one can forecast fairly directly that, even if all these people were involved in engineering and science, there will not be enough of them to meet future demand. I think the problem is in fact worse than this because Computer Science is still a young and maturing discipline.

So another interest I have in seeing Microsoft engage with the scientific community is in helping to bridge the divide between the Computer Science community and the broader world of research and engineering. My personal belief is that what we currently know as computing is going to have to evolve substantially—and what we know as programming is going to have to evolve even more dramatically. Every person who is involved in software development will struggle to deal with the complexity that comes from assembling ever larger and more complicated and interconnected pieces of software. Microsoft, as a company that aspires to be the world leader in providing software tools and platforms, is thinking deeply about how to solve those problems. One of the features that attracts me toward the world of high performance computing is that it is a world made up of people who have daily problems that need to get done, who live in an engineering environment but who are frequently at the bleeding edge in terms of the tools and techniques. And frankly there is a level of aggressiveness in this community that cannot really exist in basic business IT operations, particularly not at the scale where people are attempting to solve big new problems. So for all these reasons, Bill Gates and I decided that even though technical computing is not going to be the world's largest software market, this is a strategic market in the sense that the HPC community can help us all better understand these challenging problems. We therefore hope that together we can help move the ball forward in some of these very difficult areas. As we look downstream and contemplate some fairly radical changes in the nature of computing itself and the need for software tools to deal with that, we also expect that this community is a place from which technical leaders can emerge. We would like to be a part of that.

We think that Microsoft has some assets that could really make a difference for the growing community of people who will need to adopt HPC technologies for their business or their research. Before too long, these people will not only want to solve the problem but will also want to be able to configure and manage these HPC systems for themselves. One thing that Microsoft can do really well is to provide good tools not only for programming but also for administration, management and security.

The Next Decade in HPC

Moore's Law and the need for new algorithms

Another area where I think we can make a difference is to explore how some of our research on algorithms in a number of different areas of Computer Science could be used in other areas of science. The breakthroughs that can come from radical concepts in the algorithmic space can be quite dramatic. Several years ago some of the most passionate researchers working to develop a vaccine for AIDS approached some researchers at Microsoft Research. They showed our researchers the algorithms that their community had been using to work with the genomic data over the past six years. They were frustrated with the level of progress that had been made and wanted our researchers' opinions about how they might better work with the data to progress more quickly in producing a vaccine.

The Microsoft Research scientists, whose areas of expertise include machine learning and machine vision, studied the algorithmic concepts the group was using. They found the algorithms were sound and said that they might be able to show the group how to make the algorithms work a little better. But they also said that a whole new class of algorithms had recently been developed to work in the realm of searching in high-dimensional spaces. Not only did the scientists help the group implement the newer algorithms, but because they also were part of a group within Microsoft Research that researches the building of powerful visualization tools, they helped the group develop a suite of visualization tools that they could use along with the algorithms. The AIDS research team was able to reprocess in six weeks the same gene data that had previously taken them six years to process. Without even changing the underpinnings of the hardware or anything else they were using, the group now approaches their research in whole new ways. These new algorithms could materially accelerate the development of an effective vaccine for AIDS.

Finally, there is another challenge ahead for all of us. Since the early days of supercomputers, exploitation of parallelism has been limited by Amdahl's Law, which basically says that it is the little part of the problem that is not parallelizable that will come back to haunt you. If you have a problem that is 90% parallelizable, even if you use 1000s of processors so this parallelizable part is done blindingly fast, you are still left with the 10% that runs at the speed of one processor and a maximum speed-up of 10. We are now entering a new era of silicon technology, one in which scalar processor performance will not improve significantly. Such a statement has never been true since the invention of digital computing as we know it, but it is one that is now likely to be true for the foreseeable future. This will herald a fundamental change for the whole IT industry.

The reason is as follows. Until there is some radical change in how we physically build transistors and design computers, we have now run into a brick wall. The problem is that everybody thinks Moore's Law has really been driving the performance gains we have seen over the last thirty years. But Gordon, when he defined Moore's Law, actually only talked about the fact that there would be the ability to double the number of transistors on a chip at an exponential rate. And this phenomenon is still going to continue for awhile. But this in itself was not what actually brought all of us faster computing. The thing that actually made it happen, and made a lot of other things go faster, was raising the clock rate. But raising the clock rate was only possible because we could lower the voltage. Now we cannot lower the voltage anymore because

we are down into electron volts. There is just no more room to keep shrinking the voltage. If you cannot lower voltage, you cannot raise the clock rate materially. Therefore, even though we could have lots more transistors, there will no longer be chips with higher clock rates. This has a very profound effect. Up to now, that last 10% of the code that you could not parallelize was manageable because you have been the beneficiary of orders of magnitude improvement in scalar performance. Now such easy gains are over—or they will be in your lifetime.

So there are some very big challenges that we will all have to face. I contend that one of them is that we are all going to have to think at some level about new algorithms. One aspect of the future is already becoming clear. We will need to learn how best to exploit new multi-core CPU architectures and develop tools to support software development on such architectures. This is a real challenge to the parallel computing community. To benefit from these new machines we may have to change our programming methodology in more radical ways than we have really been comfortable doing in the past. We think about this in Microsoft every single day as we look out on a 10-20 year horizon. If you are not already thinking about this you should be. If you layer this problem on top of all the other ones, then the challenges ahead for the world's Computer Science research community are amplified quite dramatically. All of these things are what led Bill Gates and I to believe that it is strategically important, no matter what your field of software expertise, to know and think deeply about high performance computing. It is this community that will be at the forefront of examining these hard problems and in finding solutions.

PERSPECTIVES

PITAC's Look at Computational Science

In June 2004, the President's Information Technology Advisory Committee (PITAC) was charged by John Marburger, the President's Science Advisory, to respond to seven questions regarding the state of computational science. Following over a year of hearings and deliberations, the committee released its report, entitled *Computational Science: Ensuring America's Competitiveness*, in June 2005. What follows are some of my personal perspectives on computational science, shaped by the committee experience. Any wild eyed, crazy ideas should be attributed to me, not to the committee.

Based on community input and extensive discussions, the PITAC computational science report¹ included the following principal finding and recommendation.

Principal Finding. *Computational science is now indispensable to the solution of complex problems in every sector, from traditional science and engineering domains to such key areas as national security, public health, and economic innovation. Advances in computing and connectivity make it possible to develop computational models and capture and analyze unprecedented amounts of experimental and observational data to address problems previously deemed intractable or beyond imagination. Yet, despite the great opportunities and needs, universities and the Federal government have not effectively recognized the strategic significance of computational science in either their organizational structures or their research and educational planning. These inadequacies compromise U.S. scientific leadership, economic competitiveness, and national security.*

Succinctly, the principal finding highlights the emergence of computational science as the third pillar of scientific discovery, as a complement to theory and experiment. It also highlights the critical importance of computational science to innovation, security and scientific discovery, together with our failure to embrace computational science as a strategic, rather than a tactical capability. In many ways, computational science has been everyone's "second priority," rather than the unifying capability it could be.

Principal Recommendation. *Universities and the Federal government's R&D agencies must make coordinated, fundamental, structural changes that affirm the integral role of computational science in addressing the 21st century's most important problems, which are predominantly multidisciplinary, multi-agency, multisector, and collaborative. To initiate the required transformation, the Federal government, in partnership with academia and industry, must also create and execute a multi-decade roadmap directing coordinated advances in computational science and its applications in science and engineering disciplines.*

The principal recommendation emphasizes the silos and stovepipes (choose your favorite analogy) that separate disciplinary domains within computational science. There was widespread consensus from both those who testified and those on the committee that solving many of the most important problems of the 21st century will require integration of skills from diverse groups. The group also felt deeply that current organizational structures in academia and government placed limits on interdisciplinary education and research.

Based on this recognition, the committee's principal recommendation was to create a *long-term, regularly updated strategic roadmap* of technologies (i.e., software, data management, architectures and systems, and programming and tools), application needs and their interplay. The long term, strategic aspect of this recommendation cannot be over-estimated. Many of our

Dan Reed

Renaissance Computing Institute,
University of North Carolina at Chapel Hill

¹ The PITAC report on computational science can be downloaded from www.nitrd.gov. Paper copies of the report can be requested there as well.

most important computational science challenges cannot be solved in 1-3 years. Nor is a series of three year plans the same as a 10-15 year plan.

Substantial, sustained investment, driven by multi-agency collaboration, is the only approach that will allow us to escape from our current technology quandary—high-performance computing systems that are based on fragile software and an excessive emphasis on peak performance, rather than sustained performance on important applications. Simply put, today's computational science ecosystem is unbalanced, with a software and hardware base that is inadequate to keep pace with and support evolving application needs. By starving research in enabling software and hardware, the imbalance forces researchers to build atop crumbling and inadequate foundations. The result is greatly diminished productivity for both researchers and computing systems.

Similarly, we must embrace the data explosion from large-scale instruments and ubiquitous, microscale sensors—the personal petabyte is in sight! Given the strategic significance of this scientific trove, the Federal government must provide long-term support for computational science community data repositories. HPC cannot remain synonymous with computing, but must be defined broadly to include distributed sensors and storage.

Opportunities for the Future

In the 19th and 20th centuries, proximity to transportation systems (navigable rivers, seaports, railheads, and airports) was critical to success. Cities grew and developed around such transportation systems, providing jobs and social services. In today's information economy, high-speed networking, data archives and computing systems play a similar role, connecting intellectual talent across geographic barriers via virtual organizations (VOs)—teams drawn from multiple organizations, with diverse skills and access to wide ranging resources, that can coordinate and leverage intellectual talent. Two examples serve to illustrate both the challenges and the opportunities that could accrue from visionary application of computational science.

Disaster Response. Hurricane Katrina drove home the centrality of VOs. In computational science terms, a rapid response VO would include integrated hurricane, storm surge, tornado spawning, environmental, transportation, communication and human dynamics models, together with the experts needed to analyze model outputs and shape public policy for evacuation, remediation and recovery. Computationally, solving such a complex problem requires real-time data fusion from wide arrays of distributed sensors, large and small; coupled, computational intense environmental models; and social behavior models. There are thousands of such 21st century problems, each awaiting application of computational science tools and techniques.

Systems Biology. The fusion of knowledge from genomics, protein structure, enzyme function and pathway and regulatory models to create systemic models of organelles, cells and organisms and their relation to the environment is one of the great biological challenges of the 21st century. By combining information from experiments, data gleaned from mining large-scale archives (e.g., genomic, proteomic, structural and other data), and large-scale biological simulations and computational models, we can gain insights into function and behavior—understanding life in a deep way. The time is near to mount a multidisciplinary effort to create artificial life, a computational counterpart to Craig Venter's minimal genome project. Such an effort would combine engineering, genomics, proteomics and systems biology expertise, with profound implications for medicine and deep insights into biology.

The computational science opportunities have never been greater. It is time to act with vision and sustained commitment.

CTWatch QUARTERLY

ISSN 1555-9874

Volume 1 Number 4 November 2005

PUBLISHERS

Fran Berman, Director of SDSC
Thom Dunning, Director of NCSA

EDITOR-IN-CHIEF

Jack Dongarra, UTK/ORNL

MANAGING EDITOR

Terry Moore, UTK

EDITORIAL BOARD

Phil Andrews, SDSC
Andrew Chien, UCSD
Tom DeFanti, UIC
Jack Dongarra, UTK/ORNL
Jim Gray, MS
Satoshi Matsuoka, TiTech
Radha Nandkumar, NCSA
Phil Papadopoulos, SDSC
Rob Pennington, NCSA
Dan Reed, UNC
Larry Smarr, UCSD
Rick Stevens, ANL
John Towns, NCSA

CENTER SUPPORT

Greg Lund, SDSC
Karen Green, NCSA

PRODUCTION EDITOR

Scott Wells, UTK

GRAPHIC DESIGNER

David Rogers, UTK

E-INFRASTRUCTURE: EUROPE MEETS THE E-SCIENCE CHALLENGE

GUEST EDITORS TONY HEY AND ANNE TREFETHEN

AVAILABLE ON-LINE AT:
<http://www.ctwatch.org/quarterly/>



E-MAIL CTWatch QUARTERLY:
quarterly@ctwatch.org



<http://icl.cs.utk.edu/>



<http://www.ncsa.uiuc.edu/>

SDSC

<http://www.sdsc.edu/>

© 2005 NCSA/University of Illinois Board of Trustees
© 2005 The Regents of the University of California



CTWatch Quarterly is a publication of the CyberInfrastructure Partnership (CIP). <http://www.ci-partnership.org/>