Semantic Web and Grid Computing

Carole GOBLE
University of Manchester, UK

David DE ROURE
University of Southampton, UK

Abstract. Grid computing involves the cooperative use of geographically distributed resources, traditionally forming a 'virtual supercomputer' for use in advanced science and engineering research. The field has now evolved to a broader definition involving flexible, secure, coordinated resource sharing among dynamic collections of individuals, institutions and resources. This is closely related to the Semantic Web vision. In this chapter we introduce grid computing and discuss its relationship to the Semantic Web, explaining how grid applications can and should be applications of the Semantic Web; this is illustrated by a case study drawn from the life sciences. We indicate how Semantic Web technologies can be applied to grid computing, we outline some e-Science projects using Semantic Web technologies and finally we suggest how the Semantic Web stands to benefit from grid computing.

1. Introduction

In the mid 1990s Foster and Kesselman proposed a distributed computing infrastructure for advanced science and engineering, dubbed 'The Grid' [1]. The name arose from an analogy with an electricity power grid: computing and data resources would be delivered over the Internet seamlessly, transparently and dynamically as and when needed, just like electricity. The Grid was distinguished from conventional distributed computing by a focus on large-scale resource sharing, innovative science-based applications and a high performance orientation. In recent years the focus has shifted away from the high performance aspect towards a definition of the 'Grid problem' as "flexible, secure, coordinated resource sharing among dynamic collections of individuals, institutions, and resources – what we refer to as virtual organizations." [2]

The Semantic Web Activity statement of the World Wide Web Consortium (W3C) describes the Semantic Web as "…an extension of the current Web in which information is given well-defined meaning, better enabling computers and people to work in cooperation. It is the idea of having data on the Web defined and linked in a way that it can be used for more effective discovery, automation, integration, and reuse across various applications. The Web can reach its full potential if it becomes a place where data can be shared and processed by automated tools as well as by people." [3]

The Grid is frequently heralded as the next generation of the Internet. The Semantic Web is proposed as the (or at least a) future of the Web [4]. Although until very recently the communities were orthogonal; the visions are not, and neither should be the technologies. Grid computing applications can and should be seen as Semantic Web applications [5].

In this chapter we provide an overview of grid computing and discuss its relationship to the Semantic Web. We commence, in sections 2 and 3, with an introduction to the origins and evolution of grid computing. In section 4 we discuss the relationship between the Grid and the Semantic Web visions, and then focus on a life sciences grid computing scenario in section 5. After a recap of Semantic Web technologies in section 6, we look in section 7 at the ways in which such a grid computing scenario can benefit from the Semantic Web. In section 8 we introduce some e-Science projects which are using Semantic Web technologies, and in the closing discussion of section 9 we suggest how the Semantic Web stands to gain from grid computing.

2. Origins of Grid Computing

The origins of the Grid lay in 'metacomputing' projects of the early 1990s, which set out to build virtual supercomputers using networked computer systems – hence the early emphasis on high performance applications. For example, the I-WAY project [6] was a means of unifying the resources of large US supercomputing centres, bringing together high performance computers and advanced visualization environments over seventeen sites. In contrast, the FAFNER (Factoring via Network-Enabled Recursion) project ran over networked workstations – described as 'world-wide distributed computing based on computationally enhanced Web servers' [7]. In both cases the goal was computational power and the challenge was finding effective and efficient techniques to utilise the networked computational resources, be they supercomputers or workstations.

Increasing the computational power by combining increasing numbers of geographically diverse systems raises issues of heterogeneity and scalability. These distributed computing infrastructures involve large numbers of resources – both computational and data – that are inevitably heterogeneous in nature and might also span numerous administrative domains. Scalability brings a number of challenges: the inevitability of failure of components, the significance of network latency so that it is necessary to exploit the locality of resources, and the increasing number of organisational boundaries, emphasising authentication and trust issues. Larger scale applications may also result from the composition of other applications, which increases the complexity of systems.

Rather than developing a series of 'vertical' grid applications, the vision of the Grid is an infrastructure which delivers computing and data resources seamlessly, transparently and dynamically as and when needed. This involves the development of middleware to provide a standard set of interfaces to the underlying resources, addressing the problems of heterogeneity. The Globus project [8], which has origins in I-WAY, has developed the best established grid middleware in current use. The Java-based UNICORE (UNiform Interface to COmputing REsources) project has similar goals [9].

The Grid priorities largely reflected the community that proposed it, that of High Energy Physics. Planned large-scale experiments, such as the Large Hadron Collider (LHC), capture and filter petabytes of data in a few seconds and complex simulations take months of computational processing. Subsequently, the benefits of grid computing have become apparent across a range of disciplines, such as life sciences.

Major exemplars of 'traditional' Grid include the following projects:

- The Information Power Grid (IPG) Project [10] is NASA's high performance computational grid that set out to establish a prototype production Grid environment. It has proven to be a significant Grid deployment, with a service-oriented approach to the architecture.
- The European DataGrid project [11] is setting up a computational and data-intensive Grid of resources for the analysis of data coming from scientific exploration such as LHC. It is led by CERN and funded by the European Union.

- The International Virtual-Data Grid Laboratory (iVDGL) for Data Intensive Science [12] has undertaken a very large-scale international deployment to serve physics and astronomy, building on the results of projects like DataGrid.
- TeraGrid aims to deploy ‘the world's largest, fastest, most comprehensive, distributed infrastructure for open scientific research’ [13]. It is based on Linux Clusters at four TeraGrid sites, with hundreds of terabytes of data storage and high-resolution visualisation environments, integrated over multi-gigabit networks.

The provision of computational resources in support of grid applications is supplemented by support for human interaction across the grid, known as Access Grid (AG) [14], which is designed to support group to group communication such as large-scale distributed meetings, collaborative work sessions, seminars, lectures, tutorials and training. Access Grid nodes are dedicated facilities that explicitly contain the high quality audio and video technology necessary to provide an effective user experience; they also provide a platform for the development of visualisation tools and collaborative work in distributed environments, with interfaces to grid software.

Given the nature of the Grid, there is clearly a role for a standardisation effort to facilitate interoperability of grid components and services, and this is provided by the Global Grid Forum (GGF). This is a community-initiated forum of individuals working on grid technologies, including researchers and practitioners. GGF focuses on the development and documentation of ‘best practices’, implementation guidelines and standards with ‘an emphasis on rough consensus and running code’, and has operated a series of international workshops [15].

3. Evolution of the Grid

Although motivated by a focus on high performance computing for High Energy Physics, the Grid approach is clearly applicable across a broad spectrum of scientific and engineering applications which stand to benefit from the integration of large scale networked resources. There is considerable investment in grid computing in the US, Europe and throughout the world. As further applications have been explored, the Grid has evolved in two dimensions both highly relevant for the Semantic Web: architecture and scope. These are explored in this section.

3.1 Architectural evolution: the service-based Grid

In order to engineer new grid applications it is desirable to be able to reuse existing components and information resources, and to assemble and co-ordinate these components in a flexible manner. The requirement for flexible, dynamic assembly of components is well researched in the software agents community [16] and is also addressed by the Web Services model, which has become established since the first ‘Simple Object Access Protocol’ (SOAP) standard was proposed in 1998.

The creation of Web Services standards is an industry-led initiative, with some of the emerging standards in various stages of progress through the W3C [17]. The established (sometimes de facto) standards, built on the Web languages XML and XML Schema as a transport mechanism, form layers to separate the concerns of transfer, description and discovery. Messages between services are encapsulated using SOAP; services are described using the Web Services Description Language (WSDL); services are registered for publication, finding and binding using Universal Description Discovery and Integration (UDDI).

The increasing acceptance of a service-oriented approach has led to a new service-oriented vision for the Grid: the Open Grid Services Architecture (OGSA) [18]. This brings the Grid in line with recent commercial and vendor approaches to loosely coupled middleware. Consequently, the e-Science and e-Commerce communities can benefit from each other, using industrial-strength tools and environments from major vendors.

However, the Grid's requirements mean that Grid Services considerably extend Web Services. Grid service configurations are highly dynamic and volatile, large and potentially long-lived. A consortium of services (databases, sensors and compute resources) undertaking a complex analysis may be switching between sensors and computers as they become available or cease to be available; hundreds of services could be orchestrated at any time; the analysis could be executed over months. Consequently, whereas Web Services are persistent (assumed to be available) and stateless, Grid Services are transient and stateful. Different priorities are also given to issues such as security, fault tolerance and performance. The influence of Grid Services has led, for example, to extensions in WSDL to deal with service instances and their state.

To achieve the flexible assembly of grid components and resources requires not just a service-oriented model but information about the functionality, availability and interfaces of the various components, and this information must have an agreed interpretation that can be processed by machine. Hence the emphasis is on service discovery through metadata descriptions, and service composition controlled and supported by metadata descriptions. Metadata has become key to achieving the Grid Services vision.

3.2 Scope evolution: the Information/Knowledge Grid

While the service-oriented view emerged to address the 'grid problem', another movement broadened the view of the Grid. Many e-Science activities (perhaps most) are more focused on the management and interoperation of heterogeneous information.

For example, the Life Sciences community is globally distributed and highly fragmented, so that different communities act autonomously producing tools and data repositories that are built as isolated and independent systems. Few centralised repositories exist except for critical resources. Most biological knowledge resides in a large number of modestly sized heterogeneous and distributed resources, including published biological literature (increasingly in electronic form) and specialised databases curated by a small number of experts. The complex questions and analyses posed by biologists cross the artificial boundaries set by these information-generating services.

We use "information generating services" rather than databases knowingly. Information is held in databases (and thus generated from them) but is also generated by instruments, sensors, people, computational analysis and so forth. The pressing need is to *weave together* information by finding it and linking it meaningfully. Astronomy, biodiversity, oceanography, geology are all characterised by the need to manage, share, find and link large quantities of diverse, distributed, heterogeneous and changeable information.

Keith Jeffery proposed organising conceptual services into three layers, illustrated in figure 1:

- A *data/computational grid* forms the fabric of the Grid to provide raw computing power, high speed bandwidth and associated data storage in a secure and auditable way. Diverse resources are represented as a single 'metacomputer' (virtual computer), so the way that computational resources are allocated, scheduled and executed, and the way that data is shipped between processing resources, is handled here.
- An *information grid* provides homogeneous access to heterogeneous distributed information by dealing with the way that all forms of information are represented,

stored, accessed, shared and maintained. This layer orchestrates data and applications to satisfy the request, including toolkits for composing workflows, accessing metadata, visualisation, data management, and instrumentation management. The Web, and other well-known and current middleware technologies are incorporated into one framework.

- A *knowledge grid* using knowledge based methodologies and technologies for responding to high-level questions and finding the appropriate processes to deliver answers in the required form. This last layer includes data mining, machine learning, simulations, ontologies, intelligent portals, workflow reasoning and Problem Solving Environments (PSEs) for supporting the way knowledge is acquired, used, retrieved, published and maintained. A knowledge grid should provide intelligent guidance for decision makers (from control room to strategic thinkers) and hypothesis generation.

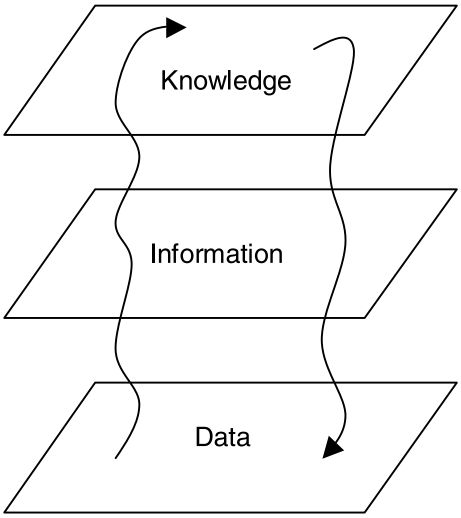

- Data mining, visualisation, simulation, problem solving methods/environments, semantic portals …
- Metadata, middleware, intelligent searching, information modelling, warehousing, workflow…
- Warehousing, distributed databases, streaming, near-line storage, large objects, access mechanisms, data staging…

Figure 1: Three conceptual layers for the Grid (Jeffery)

Each layer represents a view for, or a context of, the previous layer. Multiple interpretations are possible at the junction between each layer. Each interpretation carries the context of whom or what is viewing the data, with what prior knowledge, when, why and how the data was obtained, how trustworthy it is etc. We can imagine a frame moving from bottom to top, so each will be re-interpreted as data for the next phase. Data could be measurements, the information a collection of experimental results and the knowledge an understanding of the experiment's results or its application in subsequent problem solving.

The layered model has proved useful to promote an expansion of the kind of services a Grid should support, although it has caused some confusion. In the original proposal the Knowledge Grid was where knowledge is generated rather than held; the Information Grid is where the knowledge is encoded. This has led to others merging the Knowledge and Information Grid into one. Whatever the semiotic arguments, in this expansion of the Grid vision, metadata is clearly apparent as an essential means of filtering, finding, representing, recording, brokering, annotating and linking information. This information must be shared and must be computationally consumable.

4. Relationship between the Semantic Web and Grid Computing

We have suggested that grid applications can be seen as Semantic Web applications, a step towards the 'Semantic Grid' [5]. Figure 2, which is based on a diagram by Norman Paton, captures the relationship between the two visions. The traditional grid infrastructure extends the Web with computational facilities, while the Semantic Web extends it with richer semantics. Hence we suggest that the evolving Grid falls further up the 'richer semantics' axis, as indicated by the dotted line in the figure.

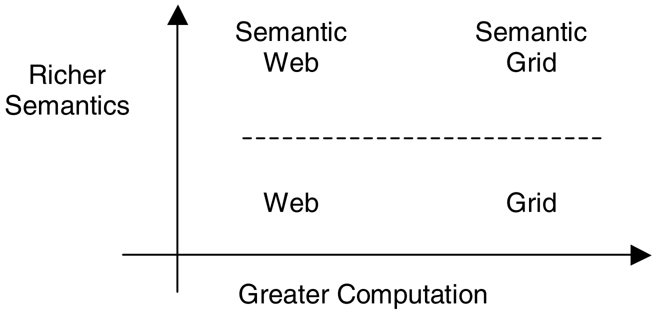

Figure 2: The Semantic Web and the Grid

Computationally accessible metadata is at the heart of the Semantic Web. The purpose of the Semantic Web is to describe a resource (anything with a URL) with what it is *about* and what it is *for*. Metadata turns out to be the fuel that powers engines that drive the Grid. Even before the Grid Service movement, metadata lay at the heart of the architecture diagrams of many grid projects. Figure 3 illustrates such an architecture.

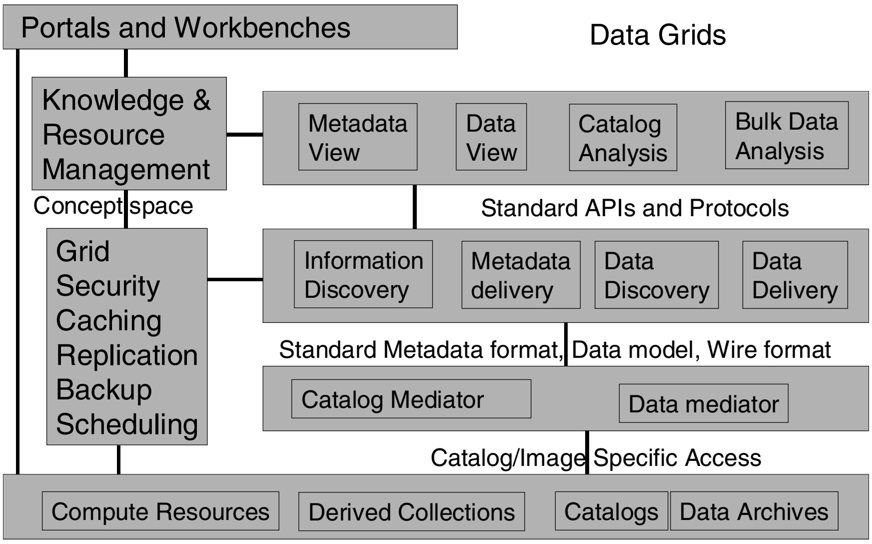

Figure 3: Example of Grid Architectures demonstrating the prevalence of metadata (NPACI)

The architectural/scope dimensions along which the Grid has evolved are orthogonal. We can use a similar duality when discussing the 'Semantic Grid' [5]. We can distinguish between a *Grid using semantics* in order to manage and execute its architectural components (a Semantic Grid Services perspective) and a *Grid of semantics* based on knowledge generated by using the Grid – semantics as a means to an end and also as an end itself. The distinction is fuzzy of course and metadata will have a dual role. In this chapter we focus on the realisation of a Semantic Grid as a grid that uses Semantic Web technologies as appropriate, throughout the middleware and application.

To achieve the full richness of the e-Science vision – the 'high degree of easy-to-use and seamless automation and in which there are flexible collaborations and computations on a global scale' [5] – also requires the richness of the Semantic Web vision. This may include, for example, distributed inference capabilities, and working with inconsistent and changing data, metadata and ontologies. This is the territory above the dotted line in figure 2, and for practitioners it is important to distinguish between what is possible now and what may be possible in the future.

5. A Life e-Science Scenario

The Grid has been driven by e-Science. In this section we use an e-Science scenario from the Life Sciences to illustrate the full extent of the Grid vision, and draw comparisons with the Semantic Web.

Large-scale science is increasingly carried out through distributed global collaborations that will require access to very large data collections, very large scale computing resources and high performance visualisation. In practice, biology has already moved to large interdisciplinary teams distributed throughout the world working together on specific problems; e.g. the Human Genome Project.

The broad range of computational grid applications in this field includes protein-folding simulations, large scale sequence pattern matching, gene expression microarray mining and combinatorial chemistry. The computational power needed to model metabolic pathways or cells is huge. However, equally pressing is the fact that post-genomics and high throughput experimentation is promising to overwhelm the community with an avalanche of data that needs to be organised and harnessed. Advances in experimental techniques enable the rapid production of large volumes of data. The introduction of DNA microarray[1] technology is a good example of this.

Biological data is often complex, represented by different media, variable in quality, stored in many places, difficult to analyse, frequently changing and mostly comprised of incomplete data sets. Analysis methods to handle the different types of data are constantly and rapidly evolving. The questions asked of the data, and the computational analyses to ask them, are more complicated: multiple species rather than single species; whole genome rather than single gene; whole metabolic lifecycle rather than single biological process. Consequently, the traditional scientific experimental methods are supplemented with 'in silico experiments', for example, the prediction of genes and the metabolic pathways they encode from the genomic DNA of an organism.

Consider a biologist in a team examining the effect of neurotransmitters on circadian rhythms in Drosophila. Before conducting a microarray experiment she checks the literature and the laboratory's online lab books for whether any other similar experiment has taken place and if the data was already available. A sample is logged into a database and labelled. A set of parameters for the machine are inferred by past experiences of similar data for similar experiments, both by others and by the biologist from those used on previous runs (the microarray machine recognised the scientist from the log). The

---

[1] A microarray, in some ways resembling a computer chip, contains thousands of spots of known DNA samples. Hence in a single experiment, thousands of genetic elements can be studied simultaneously.

parameters are recorded with the output results, which are stored in her personal database alongside the image results. The results are immediately accessible by her from her office where she analyses them with a number of specialist statistical computations and a complex interactive time-series visualisation, both of which dynamically exploit a number of available computational resources to get better performance. The visualisation is examined collaboratively with a colleague on a remote site. Both scientists attach online personal notes to the results they share between themselves but are otherwise private.

Several proteins look interesting. In order to find the way that the functions of the clusters of proteins interrelate, three databases are linked together to find the proteins, group them into families and group and visualise them by their functions, as described by a controlled vocabulary held in a third database. Huge data repositories and complex queries will demand that the computational load is spread over the most available high performance machines. Other services will have to resolve the semantic differences between database schemas and content in different resources.

Papers, in free text, quoted in the database entries and extracted online from the Medline digital library reveal that, in certain circumstances, it could control genes related to the gene of interest. The system recommends other scientists who have published work or experiments that are related. From this she discovers the gene is a transcription factor. Valuable information may have to be extracted from free format text fields of flat files. These services require additional program logic to resolve problems associated with semantic heterogeneity of life science data sets.

The system inspects the biologist's laboratory's various 'transcriptome' databases, and discovers that genes that were co-regulated with the original gene also share a target site. This discovery activity requires a number of specialised services to interact with one another, some of which are data processing while others are computationally expensive. This information is added to a public database with a link to the workflow of database interrogations and analysis tools that lead to the discovery, including versions of databases, parameter settings, versions of the algorithms and the lab that made the discovery.

Suppose the geneticist wants to repeat the discovery process for all other over-expressed genes in the microarray data set. This is a data pipelining activity used to prove hypotheses, and requires the same chain of services to be applied to different data sets. If the scientist wants to express the methodology for this process in published work, she will want to preserve as much information about how services were used so her colleagues can replicate the activity. If the geneticist is working in a pharmaceutical company and the research proves commercially valuable, the company may want the ability to trace which proprietary was were used to support a conclusion. Both of these scenarios require the Grid to have an aspect of provenance, in order to track state information about various stages of work in a research activity.

Other scientists with appropriate access rights to this database who have run an analysis that included the gene in the last month are automatically notified with this new information. Another scientist incorporates the results into a simulation of a metabolic pathway they are running, using a problem-solving environment. The simulation is monitored by various colleagues around the world, who record both private and public observations. The simulation and its results are added to a public database, and trigger new simulations automatically.

Compare this scenario with that proposed as a vision for the Semantic Web by Tim Berners-Lee in the seminal Scientific American article [4], and those proposed by others [19, 20]. The similarities are striking. The Grid provides, and the scenario demands, a more comprehensive infrastructure, encompassing computational processes, security, authentication, accounting and so forth. However, a Grid application can be thought of as a Semantic Web application.

6. A Refresher on the Semantic Web Technologies

With the above scenario in mind, here we recap the vision of the Semantic Web. It is to evolve the web into one where information and services are understandable and useable by computers as well as humans. Automated processing of web content requires explicit machine-processable semantics associated with those web resources.

To realise this vision is in some ways mundane and in others daunting, depending on your ambition. As McBride points out [21], simple metadata and simple queries give a small but not insignificant improvement in information integration. Others have more ambitious ideas of an environment where software agents are able to discover, interrogate and interoperate resources dynamically, building and disbanding virtual problem solving environments [4], discovering new facts, and performing sophisticated tasks on behalf of humans. The key point is to move from a web where semantics are embedded in hard-wired applications to one where semantics are explicit and available for automated inference.

The core technologies proposed for the Semantic Web have their roots in distributed systems and information management:

- unique identity of resources by a URI and namespace scheme;
- annotation of resources with metadata for subsequent querying or manipulation;
- shared ontologies to supply the terms used by metadata in order that the applications or people that use it share a common language and a common understanding of what the terms mean (their semantics);
- inference over the metadata and ontologies such that unasserted facts or knowledge are inferred.

The minimal components needed for this vision include annotation mechanisms, repositories for annotations and ontologies with associated query and lifecycle management, and inference engines that are resilient, reliable and perform well. Then we need the tools to acquire metadata and ontologies (manually and automatically), describe resources with metadata and express metadata using ontologies, and for versioning, update, security, view management and so on. Such an environment forms a sandpit for search engines, information brokers and ultimately the 'intelligent' agents referred to by [4], that tend it, harvest it and enable it to blossom.

The 'layer cake' model of Figure 4 depicts Tim Berners-Lee's Semantic Web vision as a set of technology layers. URIs and Unicode provide standard ways to define references to entities and to exchange symbols. XML and XMLS enable machine-processable information. These provide the syntactic underpinnings of the Semantic Web.

The first layer of the Semantic Web itself is RDF which provides a means to represent the metadata that is needed to describe any kind resource, from a web page to a web service; "a foundation for processing metadata; it provides interoperability between applications that exchange machine-understandable information on the Web" [3]. RDF Schema defines a simple modelling language on top of RDF. Ontologies enable software agents to agree on the meaning of the content within web resources, and by providing definitions of key terms, allow them to interpret the meaning of metadata attached to resources. These definitions will be written using formal (logical) languages that facilitate automated reasoning. RDFS itself has expressive limitations; DAML+OIL [22], which is the basis of the OWL Web Ontology Language being designed by the W3C Web Ontology Working Group, provides the means to represent such ontologies, extending RDF Schema to define terminology in a restricted subset of first order logic. The automated reasoning supported by DAML+OIL/OWL, and further rule languages such as RuleML [23], infers new metadata and knowledge to classify services, discover alternative services or resources and ensure interoperability between services.

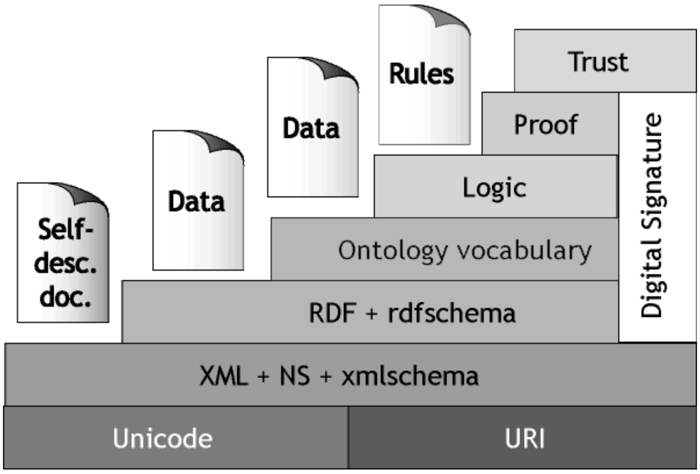

Figure 4. The Semantic Web Layer Cake (Berners-Lee, from XML2000 address)

The proof and trust layers are the least defined of the Semantic Web, but are intended to support provenance and confidence in the results. Proof is effectively exposing the reasoning used during inference combined with whether the facts and knowledge used to infer it are trustworthy. The confirmation that resources, metadata, knowledge and proofs are genuine is through digital signatures. These levels also relate to security and privacy. [24] suggests that these are applications where the other layers are languages.

7. Semantic Web Technologies in the Grid

Using the earlier scenario, we can explore some relationships with the Semantic Web vision and the opportunities for Semantic Web technologies.

7.1 Metadata based middleware

Metadata appears throughout the scenario and throughout a Grid application and all levels of the Grid. Consequently the metadata technologies developed for the Semantic Web are applicable and relevant. The following examples are generic to e-Science and hence pertain to the middleware rather than specific application domains:

- *Annotations* of results, workflows and database entries could be represented by RDF graphs using controlled vocabularies described in RDF Schema and DAML+OIL;
- *Personal notes* can be XML documents annotated with metadata or RDF graphs linked to results or experimental plans;
- Exporting *results* as RDF makes them available to be reasoned over;
- RDF graphs can be the “glue” that associates all the components (literature, notes, code, databases, intermediate results, sketches, images, workflows, the person doing the experiment, the lab they are in, the final paper) of an experiment, both in its *in silico* and ‘at the bench’ parts;
- The *provenance* trails that keep a record of how a collection of services were orchestrated so they can be replicated or replayed, or act as evidence, could for example be Web Service Flow Language (WSFL) [25] logs annotated by RDF-based metadata;
- Personalisation – the sculpting of resources into a personal “knowledge landscape” – is facilitated by descriptions relating to people and tasks.

Metadata is also exercised within the machinery of the grid computing infrastructure, for example:

- At the data/computation layer: classification of computational and data resources, performance metrics, job control, management of physical and logical resources;
- At the information layer: schema integration, workflow descriptions, provenance trails;
- At the knowledge layer: problem solving selection, intelligent portals;
- Governance of the Grid, for example access rights to databases, personal profiles and security groupings;
- Charging infrastructure, computational economy, support for negotiation; e.g. through auction model.

The metadata needs an agreed representation and, like any information, it has a lifecycle: it must be created, published, maintained and, ultimately, discarded. The creation of agreed metadata schema may be a community-led process. For example, bioinformatics content may be described using practices established in that community, while the grid infrastructure metadata schemas may be led by an organisation such as the Global Grid Forum. In the future, grid middleware could come with its own metadata schema and also support those of the application and users.

7.2 Dynamic marshalling and combining of resources

Complex questions posed by biologists require the fusion of evidence from different, independently developed and heterogeneous resources. The 400 or so public data repositories in active service in biology have different formats, interfaces, structures, coverage, etc. The Web and the Data Grid guarantee a certain level of interoperability in retrieving and accessing data. The next level of interoperability is not just making data available, but understanding what the data means so that it can be linked in appropriate and insightful ways, and providing automated support for this integration process [26]. Biologists are required to orchestrate resources in broadly two ways: (a) *Workflow orchestration*: Process flows, or workflows coordinating and chaining services using a systematic plan, are the manifestation of *in silico* experiments, allowing us to capture and explicitly represent the e-Scientist's experimental process; and (b) *Database integration*: dynamic distributed query processing, or the creation of integrated databases through virtual federations (e.g. TAMBIS [27]), portals or data warehouses (e.g. GIMS [28]).

Some schema belong to the application domain but others may be more generic (horizontal) to characterise grid computing resources. For example, coordinated distributed resource sharing applies to resources that are machines. Computationally intensive data analysis and predictive modelling can take advantage of spare resources available on machines connected to the Grid. Resources are discovered, allocated and disbanded dynamically and transparently to the user. Mediation between different Grid resource brokering models such as Unicore and Globus is a similar problem to mediating between two databases.

We can use Semantic Web technologies to:

- Represent the syntactic data types of Life Science objects using XML Schema data types, and use name spaces with URIs to uniquely identify instances of Life Science objects (the Life Science Identifer or LSID) [29, 30];
- Represent domain ontologies for the semantic mediation between database schema [26], an application's inputs and outputs, and workflow work items [31];

- Represent domain ontologies and rules for parameters of machines or algorithms to reason over allowed configurations;
- Use reasoning over execution plans, workflows and other combinations of services to ensure the semantic validity of the composition [32];
- Use RDF as a common data model for merging results drawn from different resources or instruments;
- Capture the structure of messages that are exchanged between components. This is an example where RDFS itself may appear inadequate (for example, it does not support cardinality constraints) but the design of RDF Schema can be informed by a knowledge of DAML+OIL.

There will not be just one ontology, one interpretation of results nor one set of metadata per data item. Brokering the diversity of interpretations is as much a Semantic Web vision.

7.3 The descriptive nature of the information

Partially as a result of a strong document publication ethos, knowledge in biology manifests itself in the literature and in elaborate metadata or "annotations" attached to raw data. Annotations are the accumulated knowledge attributed to a sequence, structure, protein, etc and are typically semi-structured texts. Many repositories and results are exported as flat files; the range of XML formats in Life Sciences is large and accelerating.

The premise is that a scientist will read and interpret the texts, but this makes automatic processing hard and is not sustainable given the huge amount of data becoming available. Where can the Semantic Web technologies contribute? This problem is the very problem the Semantic Web is intended to address.

The Life Sciences community has a familiarity and experience with descriptive metadata and the controlled vocabularies or ontologies required to manage it, for example the Gene Ontology [33]. Data is frequently integrated not at the schema level but at the content level: resources use a shared ontology or controlled vocabulary and link through the shared values, sometimes known as domain maps [34]. The GOBO community [35] recently recommended that the plethora of bioontologies coming on stream be encoded in DAML+OIL as a common exchange language and because the reasoning support is crucial when building large collaborative community ontologies.

7.4 The computational inaccessibility of information and applications

Many bioinformatics repositories and applications have simple call interfaces without APIs or query languages. Many have interactive "point and click" visual interfaces, which are good for people and bad for automated processing. Again, making the computationally inaccessible accessible lies at the heart of the Semantic Web. The service-oriented approach to the engineering of applications will address this. Components of legacy applications can be 'wrapped' as services and incorporated into the service description schemes.

7.5 Provenance, quality, trust, proof and the experimental process

Both the results, and the way they were obtained, are high value. What the data is like and where it came from is as important as the data itself. As data collections and analytical applications evolve, keeping track of commensurate changes is difficult, so "change notification" becomes essential. Our scientist will need to rerun their experiment if data changes, or new knowledge questions the underlying premise of the analysis. In biology, data is replicated in secondary databases through annotation. This "viral migration" of information makes the tracking of data even harder. Mistakes or discredited information are propagated and difficult to eliminate.

A distributed environment on the scale of the Grid requires a number of core services built into its fabric to govern the whole scientific environment: ownership and watermarking (who owns the resource); provenance, quality, audit, versioning (where did the data come from and when); authentication, security and confidentiality (who can access the resource); personalisation and configuration (my lab book is special to me) and so on. These are clearly applications of the Proof, Trust and Digital Signatures of the Semantic Web.

7.6 The instability of science

Information and knowledge in science is heavily contextual and often opinionated. Contexts change and opinions disagree. New information may support or contradict current orthodoxy held by an individual, a team or a community, leading to a revision of beliefs. Managing, retrieving and reusing changing and evolving content is a challenge. Dealing with changing content that is also irregular, inconsistent, incomplete, contradictory, uncertain and imprecise makes the exploitation of these knowledge assets difficult. In particular, science is not "linear" – huge jumps are made in thinking, but the past cannot be wiped. It is essential to be able to recall a snapshot of the state of understanding at a point in time. This is why provenance is so important. All inferences *must* be exposed or else the scientist will not use them.

The expectation that there are multiple assertions on a resource and that inference engines over those assertions should deal with the "dirtiness" of the world is a bedrock of the Semantic Web vision. Not only does the content change, of course, but so do the ontologies and rules that we use to infer new information. When an ontology changes in line with new ideas, this does not wipe the old inferences that no longer hold (and how do we propagate those changes?). Somehow they must continue to co-exist and be accessible.

Event notification – when any entity in the Grid changes and fires further processes – is considered such an essential of e-Science that it is a core service in the Open Grid Service Architecture specification of Grid services. In our scenario, changes to database items are registered by our scientist to replay her experiment and by others who are affected by the results she publishes. Event notification schemes typically register the topics that should be monitored for change – these topics can be described using Semantic Web ontology and metadata technologies.

Scientific data has a very long lifespan. We do not throw away scientific knowledge and as a society we have a duty of guardianship. This means that the data in the Grid and the information that is being used to support the execution of the Grid persists. In addition to tending and harvesting it, we need to 'weed' the Semantic Web.

7.7 Knowledge

In the scenario, the biologist was advised of parameter settings based on information in the Grid. She was helped to choose appropriate experiments and resources, and to plan the execution of both her *in silico* and *in vitro* experiments. She used a problem-solving environment for building and monitoring her metabolic pathway simulation. The recording, and sharing, of workflows helps improve experimental practice by avoiding unnecessary replication of in silico experiments (or in vitro experiments for that matter); it also assists in setting up equipment or computational processes in appropriate ways, and helps ensure that the conclusions drawn are fully justified by the techniques used. These are all applications of, or for, the Semantic Web – personalised agents or services [20], semantic portals onto services [36], recommender systems, intelligent searching, text mining for metadata [37] and so on.

The rhetoric of the Semantic Web often presents it as a global knowledge base. The metadata on resources are the facts; the ontologies the terms; the inference engines and rule systems the reasoning tools. Our scientist might well want to pose the question "what ATPase superfamily proteins are found in mouse?" and get the answers (a) P21958 (from the Swiss-Prot database she has permission to access); (b) InterPro is a pattern database and could tell you if you had permission and paid. (c) Attwood's lab expertise is in nucleotide binding proteins (ATPase superfamily proteins are a kind of nucleotide binding protein); (d) Smith published a new paper on this in Nature Genetics two weeks ago; (e) Jones in your lab already asked this question.

7.8 Collaborative science

Knowledge about who has published a relevant paper or asked a question is one mechanism for support of collaboration between scientists. In the scenario, the visualisation is examined collaboratively with a colleague on a remote site – this is a synchronous interaction between scientists, which is another form of collaboration that may itself be facilitated by the Grid. For example, a visualisation may be computationally intensive and make use of the high bandwidth network, especially if it is interactive and collaborative.

The Access Grid [14] can support this style of interaction, where dedicated or room-based facilities are required, or else desktop solutions are possible. Potentially the collaboration need not be office-to-office but could involve scientists working within experimental laboratories or in the field. Knowledge technologies can additionally be used for personalisation of information, tracking issues through a series of meetings, recording meetings and identifying communities of practice – see section 8.4 for an example project investigating some of these issues.

7.9 Semantic Grid Services

In the section 3 we explained that today's Grid applications are in a broad range of disciplines, are service-oriented, and require metadata for discovery and utilisation of resources. The service-oriented approach can be implemented with extensions to Web Services.

At the time of writing, the current state of describing Grid Services through semantics rather than just syntax via WSDL or simple classifications is as follows: "The service description is meant to capture both interface syntax, as well as semantics …Semantics may be inferred through the names assigned the portType and serviceType elements…Concise semantics can be associated with each of these names in specification documents – and perhaps in the future through Semantic Web or other formal descriptions" [38]. This is a challenge and an opportunity for the Semantic Web community. Bringing together the Semantic Web and Web Services has already attracted attention [39,40,41,42]. Using

DAML+OIL to build service classifications more powerful than UDDI's simple hierarchies has been explored in the myGrid project [31] (see section 8.1).

The description of a service is essential for automated discovery and search, selection, (imprecise) matching, composition and interoperation, invocation, and execution monitoring. This choice depends on metadata concerned with function, cost, quality of service, geographical location, and the original publisher. Classification of services based on the functionality they provide has been widely adopted by diverse communities as an efficient way of finding suitable services. The Universal Description, Discovery, and Integration specification (UDDI) supports web service discovery by using a service classification. In Biology, the EMBOSS suite of bioinformatics applications and repositories has a coarse classification of the 200 or so tools it contains, and free text documentation for each tool; ISYS [43] and BioMOBY [29] use taxonomies for classifying services. In the pre-Services Grid, The Metadata Directory Service (MDS) [44] and Metadata Catalog (MCAT) [45] resource directory frameworks defined the properties that can be used to query a resource.

Reasoning has a role to play, not just in the creation of the ontologies used to classify services but also in the matching of services. In Condor, a structural matching mechanism was used to choose computational resources [46]. The semantic matching possible through reasoning in languages such as DAML+OIL has been explored in Matchmaker and myGrid [31]. In an architecture where the services are highly volatile, and configurations of services are constantly being disbanded and re-organised, knowing if one service is safely substitutable by another is an essential, not a luxury.

We mentioned earlier that Grid Services extend Web Services in three important ways: service instances, soft state and long-lived service orchestrations. How this will affect the way Semantic Web technologies can describe and discover Grid services is open for research.

8. Some examples of Grid Projects using Semantic Web Technologies

In 2001 the UK government launched a $180 million programme to develop and deploy Grid technology to support the challenges of 'e-Science' – the large scale science carried out through distributed global collaborations enabled by the Internet. The emphasis is on developing the grid infrastructure through projects in close collaboration with application scientists, and includes testbeds aimed at testing grid technology in a number of distinct science and engineering areas in collaboration with industry and commerce. The UK Grid vision pays particular attention to the processes by which Grid applications contribute to the creation and delivery of information and knowledge, as well as to the underlying computational resource sharing issues.

The timing of the UK e-Science programme has intercepted the second generation of service-based information-oriented Grids. Without obligations to the legacy of the earlier generation Grid technologies, many projects have been able to adopt a service-oriented approach from the outset, with due attention to information and knowledge aspects. In fact the emphasis is on the science, for which the Grid is an infrastructure, rather than on the infrastructure itself – hence some projects have adopted a more holistic view, which starts with the e-Scientist and the laboratory rather than the socket on the wall. Consequently, the full vision of e-Science is one where there is 'a high degree of easy-to-use and seamless automation and in which there are flexible collaborations and computations on a global scale' [5]. This requires more than just metadata: it requires the broader Semantic Web vision.

We now give examples of four of these e-Science projects. All four have adopted Semantic Web technologies and a service-based approach.

8.1 myGrid (www.mygrid.org.uk)

The myGrid project aims to deliver the middleware required for a personalised collaborative problem-solving environment. The focus is on data-intensive e-Biology, specifically post-genomic functional analysis, and the provision of a distributed environment that supports the in silico experimental process, though the middleware should be generic. The aim is that the e-Scientist be able to:

- Compose workflows, integrate and query databases and access to digital libraries through information extraction from texts.
- Find and adapt workflows and services developed by others, through accessing metadata based on ontologies.
- Store partial results in local data repositories, have their own view on public repositories.
- Be better informed as to the provenance and the currency of the tools and data directly relevant to their experimental space.

The Grid becomes egocentrically based around the Scientist – ‘myGrid’. The ultimate goal is to improve both the quality of information in repositories and the way repositories are used.

myGrid uses ontologies to describe the bioinformatics services it publishes and orchestrates. A suite of ontologies expressed in DAML+OIL provides: (a) sufficiently elaborate service classifications to express the domain faithfully; (b) a vocabulary for expressing service descriptions and (c) a reasoning process to manage the coherency of the classifications and the descriptions when they are *created*, and the service discovery, matching and composition when they are *deployed*. The ontology extends the DAML-S service ontology [40]. Although the work is at an early stage, it shows that DAML+OIL provides an effective language to describe the functionality of service in such a demanding arena as biology [31].

myGrid is a good exemplar of a ‘Semantic Grid’ project, in particular because it makes use of a reasoning capability. Services are sought initially by their scientific metadata (domain dependent) and use the business metadata (domain independent) to choose between services of equivalent scientific functionality. Candidate services may not necessarily be an exact match, so the Fact system [22] is used to reason over the properties of the service descriptions expressed in DAML+OIL to infer close matches that are substitutable. Services can be discovered, matched and selected both *before* the workflow is executed or *dynamically* during its execution. Services generate data that could be the input to another service or could be stored in a repository, or both. In a workflow, we need to ensure that the type of the data (for example nucleotide sequence) matches the service’s input type. The semantic *type* of the data must match: for example, a collection of enzymes is permissible as input to BLASTp as enzymes are a kind of protein and BLASTp takes sets of proteins as an input. To guide the user in choosing appropriate operations on their data and constraining which operation should sensibly follow which, it is important to have access to the semantic type of data concerned. Consequently, service descriptions cover the type of their inputs and outputs, and the type should be carried with the data for future use in new workflows. The service discovery interface of the myGrid Portal is illustrated in figure 5.

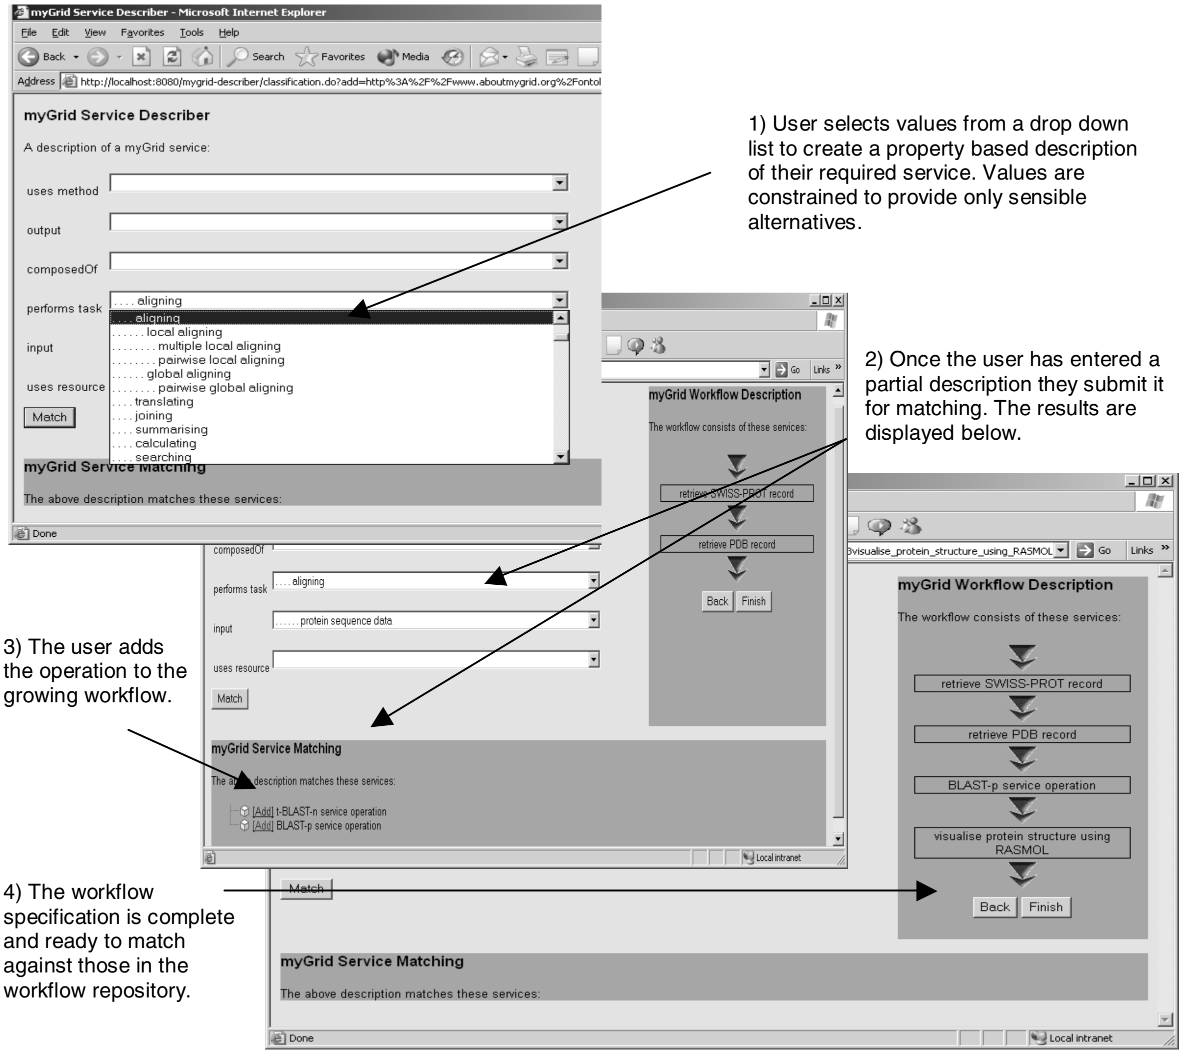

performs task . . . . aligning

input . . . . . . protein sequence data

myGrid Service Matching

The above description matches these services:

[Add] t-BLAST-n service operation
[Add] BLAST-p service operation

myGrid Workflow Description

The workflow consists of these services:

retrieve SWISS-PROT record
retrieve PDB record
BLAST-p service operation
visualise protein structure using RASMOL

1) User selects values from a drop down list to create a property based description of their required service. Values are constrained to provide only sensible alternatives.

2) Once the user has entered a partial description they submit it for matching. The results are displayed below.

3) The user adds the operation to the growing workflow.

4) The workflow specification is complete and ready to match against those in the workflow repository.

Figure 5. The service discovery interface of the myGrid Portal

8.2 Comb-e-Chem (www.combechem.org)

From an e-Science perspective, combinatorial chemistry has many similarities with the scenario in section 5. As its name suggests, combinatorial chemistry involves parallel synthesis methods that enable a large number of combinations of molecular units to be assembled rapidly – this is called a 'library'. In conjunction with this parallel synthesis there is parallel screening; e.g. for potential drug molecules. Each member of the library is tested against a target and those with the best response are selected for further study. When a significant response is found, the structure of that particular molecule is determined and used as the basis for further investigation to produce a potential drug molecule. As with the genetic databases, the size and growth rates of the databases of these molecules are dramatic.

The main objective of the Comb-*e*-Chem project is to develop an e-Science testbed that integrates existing structure and property data sources, and augments them in four ways within a grid-based resource- and knowledge-sharing infrastructure:

- by supporting new data collection, including process as well as product data, based on integration with electronic lab and e-logbook facilities;
- by integrating data generation on demand via grid-based quantum and simulation modelling to augment the experimental data;

- by developing interfaces that provide a unified view of these resources, with transparent access to data retrieval, on-line modelling, and design of experiments to populate new regions of scientific interest; and
- by providing shared, secure access to these resources in a collaborative e-science environment.

While some issues, such as provenance, are shared with myGrid, there is in Comb-e-Chem an emphasis on the experimental laboratory, on real time multimedia and collaborative aspects of e-Science, and on the design of experiments. The project has a strong automation theme, from robotics within the experimental laboratory through to automation of the e-Science workflows, and this is being addressed using the techniques of agent-based computing.

8.3 Geodise (www.geodise.org)

During the process of optimisation and design search, the modelling and analysis of engineering problems are exploited to yield improved designs. The engineer explores various design parameters that they wish to optimise and a measure of the quality of a particular design (the objective function) is computed using an appropriate model. A number of algorithms may be used to yield more information about the behaviour of a model, and to minimise/maximise the objective function, and hence improve the quality of the design. This process usually includes lengthy and repetitive calculations to obtain the value of the objective function with respect to the design variables.

The Geodise project aims to aid the engineer in the design process by making available a suite of design optimisation and search tools, Computational Fluid Dynamics (CFD) analysis packages integrated with distributed Grid-enabled computing and data resources. These resources are exposed to the user via a web-based portal. In addition, the user is guided through the design search process by an integrated knowledge base.

Ontologies serve as the conceptual backbone for knowledge sharing and management in the integrated architecture for knowledge services. The ontologies are represented in DAML+OIL and an ontology service provides a Java API giving full access to any DAML+OIL ontology available over the Internet. Annotation adds semantic content to documents or websites, thus facilitating information sharing, reuse and automatic machine processing. The OntoMat-Annotizer RDF annotation tool [47] is used to annotate workflows of optimization runs for particular design problems and then save them in a knowledge base. The semantically enriched archive can then be queried, indexed and reused later to guide future designs. A knowledge portal makes the knowledge available and accessible; provides tools for knowledge reuse and exchange; provides security infrastructure; manages knowledge resources; supports an online forum, maintains mailing lists and disseminates the latest advances of the domain.

8.4 CoAKTinG (www.aktors.org/cooakting)

The Access Grid [14], introduced in section 2, is a collection of resources that support human collaboration across the Grid, such as large-scale distributed meetings and training. The resources include multimedia display and interaction, notably through room-based videoconferencing (group-to-group), and interfaces to grid middleware and visualisation environments. It relates to the concept of a collaboratory [48], which is a distributed research centre in which scientists in several locations are able to work together.

During a meeting, there is live exchange of information, and this brings the information layer aspects to the fore. For example, events in one space can be communicated to other spaces to facilitate the meeting. At the simplest level, this might be moving through the agenda, slide transitions or remote camera control. These provide metadata, which is generated automatically by software and devices, and can be used to enrich the conference and stored for later use. New forms of information may need to be exchanged to handle the large scale of meetings, such as distributed polling and voting. Another source of live information is the notes taken by members of the meeting, including minutes and issue tracking, and the annotations that they make on existing documents. Again, these can be shared and stored to enrich the meeting. A feature of current collaboration technologies is that sub-discussions can be created easily and without intruding – these also provide enriched content.

The CoAKTinG project ('Collaborative Advanced Knowledge Technologies on the Grid') will provide tools to assist scientific collaboration by integrating intelligent meeting spaces, ontologically annotated media streams from online meetings, decision rationale and group memory capture, meeting facilitation, issue handling, planning and coordination support, constraint satisfaction, and instant messaging/presence. A scenario in which knowledge technologies are being applied to enhance collaboration is described in [49].

The combination of Semantic Web technologies with live information flows is highly relevant to grid computing and is based here on experience within the hypermedia context [50]. Metadata streams may be generated by people, by equipment or by programs (e.g. annotation, device settings, data processed in real-time) and we envisage 'metadata distribution networks'. Potentially CoAKTinG involves many ontologies, including those for the application domain, for the organisational context, for the meeting infrastructure, for devices which are capturing metadata and a constraint-based ontology for processes and products [51]. In contrast with some other projects, it requires real-time processing. For example, when someone enters the meeting, other participants can be advised immediately on how their communities of practice intersect.

9. Discussion and Conclusions

The Web was incubated by a scientific community: Physics. This community was a well-organised microcosm of the general community. It had definite and clearly articulated information dissemination needs and it had a group of smart people prepared to co-operate, and with the means and desires to do so. The state of play of the Grid today is reminiscent of the Web some years ago. At this time there is limited deployment, largely driven by enthusiasts within the scientific community (indeed, the High Energy Physics Community again), with emerging standards and a degree of commercial uptake. The same might also be said of the current state of the Semantic Web deployment, though it is not clear that the same drivers are in place as existed for Web and Grid.

Meanwhile, the Web itself has enjoyed massive deployment and continues to evolve; e.g. the shift from machine-to human communications (HTML) to machine-to-machine (XML), and the emergence of the Web Services paradigm. The requirements of one of the drivers, e-Commerce, are in line with those of e-Science. Thus the scene is set, and it is appealing to infer from these similarities that Grid and Semantic Web deployment will follow the same exponential model as the growth of the Web.

However, a typical grid application is not a typical web application. A grid application might involve large numbers of processes interacting in a coordinated fashion, while a typical Web transaction today still only involves a small number of hosts (e.g. server, cache, browser). Moreover, grid processes continually appear and disappear, while web servers persist. Achieving the desired behaviour from a large scale distributed system involves technical challenges that the Web itself has not had to address, though Web Services take us towards a similar world. This is not just a data/computation layer issue; large scale integration of resources at information level is less well developed but no less challenging.

The Semantic Web requires a metadata-enabled Web. In the same way as the components of information systems have moved to support HTML and XML in the last few years, we now need them to take on board the support for creating and maintaining metadata. Unfortunately there are many obstacles to this. In particular, manual creation of metadata is problematic. People are not always in the best position to create it and they might not produce accurate metadata, through circumstance or error; even if accurate the metadata also needs to be useful, and there is 'more than one way to describe a cat'. With grid applications we have the requirement – and also the opportunity – to automate the management of *quality* metadata. Finally, creating ontologies is hard. Addressing the problems of creating and managing ontologies is the paramount.

So there are challenges ahead for both Grid and Semantic Web. Does this mean that the proposed marriage – the Semantic Grid – is ill-fated? We argue instead that the technologies are symbiotic and the partnership is essential for both to thrive.

We have shown that the visions of the Grid and the Semantic Web are related, and that Grid applications can be Semantic Web applications. Grid computing can gain immediately from the metadata technologies of the Semantic Web, informed by the OWL Web Ontology Language. The synergy is such that the Grid will also gain from the ongoing developments of the Semantic Web – for example, from the incorporation of inference technologies in the future – taking us towards the full Semantic Grid vision.

To achieve these benefits requires that the grid computing and applications community pay due attention to the Semantic Web. This applies to vertical projects, where the Semantic Web technologies can be applied within the application domain, and also to middleware developments, which can build in the Semantic Web infrastructure. There is a cost to taking on board new technologies, and here the benefits may not always be immediate. The Semantic Web community has a role to play in supporting the initial uptake, especially as many traditional Grid developers regard themselves as systems-oriented and the adoption of knowledge technologies seems a stretch. One barrier to adoption is confusion over what can be achieved now and what is best treated as 'wait and see'.

Why should the Semantic Web researchers be interested in the Grid? It is 'just an application' but in fact a very special one – perhaps even a 'killer app' – for several reasons:

- It is a very good example of the type of application envisaged for the Semantic Web. The essence of the Grid is the power provided by large scale integration of resources, and the scale and automation of the Grid necessitates the 'universally accessible platform that allows data to be shared and processed by automated tools as well as by people'.
- It is a real application: the emphasis is on deployment and on high performance, and is on a large scale and has established communities of users. Such applications are essential to the uptake of the Semantic Web.
- The Grid genuinely needs Semantic Web technologies. Even at the most basic level, Grid developers acknowledge that 'information islands' are being created and require an interoperability solution at information level such as provided by grid middleware at data/computation level.
- It will stress Semantic Web solutions, and it raises some specific grid-related issues which will provide a useful challenge. Solutions to these issues are unlikely to be

peculiar to grid computing – related issues will surely be evident in other Semantic Web applications in the fullness of time.
- It is self-contained, with a well-defined community who already work with common tools and standards.
- Aspects of the Semantic Web could be applications of grid computing, for example in search, data mining, translation and multimedia information retrieval.

The partnership between the Semantic Web and the Grid presents an exciting vision. Each partner has obstacles to its progress, but each stands to benefit from the other. To be successful, the partnership requires disjoint communities to come together. If they do, we can look forward to the 'next generation' of the Web: one with tremendous power to enable a new paradigm in science and engineering.

Acknowledgements

This work is supported by the Engineering and Physical Sciences Research Council and Department of Trade and Industry through the UK e-Science programme, in particular the myGrid e-Science pilot (GR/R67743), the Geodise e-Science pilot (GR/R67705), the Comb-*e*-Chem e-Science pilot (GR/R67729/01), the Semantic Grid report and the CoAKTinG project (GR/R85143/01) which is associated with the 'Advanced Knowledge Technologies' Interdisciplinary Research Collaboration (GR/N15764/01). The authors wish to thank all their colleagues who are working on the 'bridge' between the Grid and the Semantic Web. Nick Jennings and Nigel Shadbolt co-authored the Semantic Grid report for the UK e-Science programme, identifying a Semantic Grid research agenda, and we are grateful to all those who have supported this initiative within the UK including Tony Hey, Mark Baker, Malcolm Atkinson, Norman Paton, Omer Rana, Keith Jeffery, Tom Rodden and members of the Architecture Task Force, and to Hans-Georg Stork for his work on the European axis. We thank Ian Foster for his help and always being ready to discuss these ideas, both Ian and Carl Kesselman for bringing grid computing to the Semantic Web community through international conferences, and the Global Grid Forum for undertaking the Semantic Grid activity. Thanks especially to Jim Hendler for encouraging and supporting the discussion of grid computing within the Semantic Web community, and to Eric Miller and Henry Thompson from W3C for their contributions across the communities. Finally thanks to everyone currently building semantic grids: to Jeremy Frey, Simon Cox and Mike Surridge for exploring semantic grids through their grid computing projects, and to all our colleagues in the myGrid, Geodise, Comb-e-Chem, CoAKTinG, GRIA and AKT project teams.

References

[1] I. Foster and C. Kesselman (eds), "The Grid: Blueprint for a New Computing Infrastructure", Morgan Kaufmann, July 1998.

[2] I. Foster, C. Kesselman, S. Tuecke "The Anatomy of the Grid: Enabling Scalable Virtual Organizations", International Journal of Supercomputer Applications, 15(3), 2001.

[3] W3C Semantic Web Activity Statement, http://www.w3.org/2001/sw/Activity/

[4] Berners-Lee,T., Hendler,J. and Lassila ,O. "The Semantic Web", Scientific American, May 2001.

[5] D. De Roure, N.R. Jenning and N.R. Shadbolt. "Research Agenda for the Semantic Grid: A Future e-Science Infrastructure". Report UKeS-2002-02 from the UK National e-Science Centre, December 2001. Also see http://www.semanticgrid.org

[6] I. Foster, J. Geisler, W. Nickless, W. Smith, S. Tuecke "Software Infrastructure for the I-WAY High Performance Distributed Computing Experiment", in Proc. 5th IEEE Symposium on High Performance Distributed Computing. pp. 562-571, 1997.

[7] FAFNER, http://www.npac.syr.edu/factoring/
[8] I. Foster and C. Kesselman, “Globus: A Metacomputing Infrastructure Toolkit”, International Journal of Supercomputer Applications, 11(2): 115-128, 1997.
[9] J. Almond and D. Snelling, “UNICORE: uniform access to supercomputing as an element of electronic commerce”, Future Generation Computer Systems, 15(1999) 539-548, NH-Elsevier.
[10] NASA Information Power Grid, http://www.ipg.nasa.gov
[11] The DataGrid project, http://eu-datagrid.web.cern.ch
[12] iVDGL - International Virtual Data Grid Laboratory, http://www.ivdgl.org
[13] TeraGrid project, http://www.teragrid.org
[14] Access Grid, http://www.accessgrid.org
[15] Global Grid Forum, http://www.gridforum.org
[16] N. R. Jennings, “An agent-based approach for building complex software systems"”, Comms. of the ACM, 44 (4) 35-41. 2001.
[17] W3C Web Services Activity, http://www.w3.org/2002/ws/
[18] I. Foster, C. Kesselman, J. Nick and S. Tuecke, “The Physiology of the Grid: Open Grid Services Architecture for Distributed Systems Integration”, presented at GGF4, Feb. 2002. See http://www.globus.og/research/papers/ogsa.pdf
[19] J. Euzenat, “Research Challenges and Perspectives of the Semantic Web”, European Commission - US National Science Foundation Strategic Research Workshop, Sophia Antipolis, France, October 2001.
[20] J. Hendler, “Agents and the Semantic Web”, IEEE Intelligent Systems Journal, March/April 2001 (Vol. 16, No. 2), pp. 30-37.
[21] B. McBride, “Four Steps Towards the Widespread Adoption of a Semantic Web”, in Proceedings of the First International Semantic Web Conference (ISWC 2002), Sardinia, Italy, June 9-12, 2002. LNCS 2342, pp 419-422.
[22] I. Horrocks, “DAML+OIL: a reason-able web ontology language”, in Proceedings of EDBT 2002, March 2002.
[23] RuleML, http:// www.dfki.uni-kl.de/ruleml/
[24] P.F. Patel-Schneider, D. Fensel, “Layering the Semantic Web: Problems and Directions”, in Proceedings of the First International Semantic Web Conference (ISWC 2002), Sardinia, Italy, June 9-12, 2002. LNCS 2342 pp 16-29.
[25] Web Services Flow Language (WSFL) Version 1.0, http://www-4.ibm.com/software/solutions/Webservices/pdf/WSFL.pdf
[26] C.A. Goble, “Supporting Web-based Biology with Ontologies”, in Proceedings of the Third IEEE International Conference on Information Technology Applications in Biomedicine (ITAB00), Arlington, VA (November 2000), pp. 384–390.
[27] C.A. Goble, R. Stevens, G Ng, S Bechofer, N. Paton, P. Baker, M. Peim and A. Brass. “Transparent access to multiple bioinformatics information sources.” IBM Systems Journal, Vol. 40, No. 2, pp 532-551, 2001.
[28] Paton, N.W., Khan, S.A., Hayes, A., Moussouni, F., Brass, A., Eilbeck, K., Goble, C.A., Hubbard, S. and Oliver, S.G., “Conceptual Modelling of Genomic Information”, Bioinformatics, Vol 16, No 6, 548-558, 2000.
[29] BioMOBY http://www.biomoby.org
[30] Life Sciences Identifier (LSID), Interoperable Informatics Infrastructure Consortium (I3C), http://www.i3c.org
[31] C. Wroe, R. Stevens, C. Goble, A. Roberts, M. Greenwood, “A suite of DAML+OIL Ontologies to Describe Bioinformatics Web Services and Data”, to appear in IJCAI Special issue on Bioinformatics Data and Data modelling.
[32] J. Cardoso and A. Sheth, “Semantic e-Workflow Composition”, Technical Report, LSDIS Lab, Computer Science, University of Georgia, July 2002.
[33] M. Ashburner et al, “Gene Ontology: tool for the unification of biology”. Nature Genetics 25: 25-29 (2000)
[34] B. Ludäscher, A. Gupta, and M.E. Martone, “Model-Based Mediation with Domain Maps”, in 17th Intl. Conference on Data Engineering (ICDE), Heidelberg, Germany, IEEE Computer Society, April 2001.
[35] Gene Ontology Consortium, http://www.geneontology.org
[36] S. Staab, J. Angele, S. Decker, M. Erdmann, A. Hotho, A. Maedche, E. Studer,, and Y. Sure, “Semantic CommunityWeb Portals”, in Proceedings of the 9th World Wide Web Conference (WWW9), Amsterdam, Netherlands, 2001.
[37] S. Staab, A. Mädche, F. Nack, S. Santini, L. Steels. “Emergent Semantics”, IEEE Intelligent Systems, Trends & Controversies, 17(1), Jan/Feb 2002, pp. 78 –86.
[38] S. Tuecke, K. Czajkowski, I. Foster, J. Frey, S. Graham and C. Kesselman, “Grid Service Specification Draft 3 7/17/2002”, http://www.gridforum.org/ogsi-wg/

[39] D. Trastour, C. Bartolini and C. Preist, “Semantic Web Support for the Business-to-Business E-Commerce Lifecycle”, in The Eleventh International World Wide Web Conference (WWW2002). pp: 89-98 2002.
[40] DAML Services Coalition, “DAML-S: Web Service Description for the Semantic Web”, in The First International Semantic Web Conference (ISWC), June, 2002, pp 348-363.
[41] D. Fensel, C. Bussler, A. Maedche, “Semantic Web Enabled Web Services”, in Proceedings of the First International Semantic Web Conference, Sardinia June 2002, Springer-Verlag LNCS 2342 pp:1-2.
[42] J. Peer, “Bringing together Semantic Web and Web Services”, in Proceedings of the First International Semantic Web Conference, Sardinia June 2002, Springer-Verlag LNCS 2342 pp: 279-291
[43] ISYS, http://www.ncgr.org/isys/
[44] Globus Monitoring and Discovery Service, http://www.globus.org/mds/
[45] Meta Information Catalog, http://www.npaci.edu/DICE/SRB/mcat.html
[46] Condor, http://www.cs.wisc.edu/condor/
[47] S. Handschuh1 and S. Staab, “Authoring and Annotation of Web Pages in CREAM”, in Proceedings of the Eleventh International World Wide Web Conference (WWW2002)
[48] V.G. Cerf et al., “National Collaboratories: Applying Information Technologies for Scientific Research”, National Academy Press: Washington, D.C., 1993.
[49] S. Buckingham Shum, D. De Roure, M. Eisenstadt, N. Shadbolt and A. Tate, “CoAKTinG: Collaborative Advanced Knowledge Technologies in the Grid”, in Proceedings of the Second Workshop on Advanced Collaborative Environments at the Eleventh IEEE Int. Symposium on High Performance Distributed Computing (HPDC-11), July 24-26, 2002, Edinburgh, Scotland.
[50] K.R. Page, D.G. Cruickshank and D. De Roure, “It’s About Time: Link Streams as Continuous Metadata”, in Proceedings of the Twelfth ACM Conference on Hypertext and Hypermedia (Hypertext '01) p.93-102. 2001.
[51] A. Tate, J. Levine, J. Dalton and A. Nixon, “Task Achieving Agents on the World Wide Web”, in Creating the Semantic Web, Fensel, D., Hendler, J., Liebermann, H. and Wahlster, W. (eds.), 2002, MIT Press.