1

Introduction

Paul Warren, Rudi Studer and John Davies

1.1. SEMANTIC WEB TECHNOLOGIES

That we need a new approach to managing information is beyond doubt. The technological developments of the last few decades, including the development of the World Wide Web, have provided each of us with access to far more information than we can comprehend or manage effectively. A Gartner study (Morello, 2005) found that 'the average knowledge worker in a Fortune 1000 company sends and receives 178 messages daily', whilst an academic study has shown that the volume of information in the public Web tripled between 2000 and 2003 (Lyman *et al.*, 2005). We urgently need techniques to help us make sense of all this; to find what we need to know and filter out the rest; to extract and summarise what is important, and help us understand the relationships between it. Peter Drucker has pointed out that knowledge worker productivity is the biggest challenge facing organisations (Drucker, 1999). This is not surprising when we consider the increasing proportion of knowledge workers in the developing world. Knowledge management has been the focus of considerable attention in recent years, as comprehensively reviewed in (Holsapple, 2002). Tools which can significantly help knowledge workers achieve increased effectiveness will be tremendously valuable in the organisation.

At the same time, integration is a key challenge for IT managers. The costs of integration, both within an organisation and with external trading partners, are a significant component of the IT budget. Charlesworth (2005) points out that information integration is needed to 'reach a better understanding of the business through its data', that is to achieve a

Semantic Web Technologies: Trends and Research in Ontology-based Systems
John Davies, Rudi Studer, Paul Warren

common view of all the data and understand their relationships. He describes application integration, on the other hand, as being concerned with sharing ‘data, information and business and processing logic between disparate applications’. This is driven in part by the need to integrate new technology with legacy systems, and to integrate technology from different suppliers. It has given rise to the concept of the service oriented architecture (SOA), where business functions are provided as loosely coupled services. This approach provides for more flexible loose coupling of resources than in traditional system architecture, and encourages reuse. Web services are a natural, but not essential, way of implementing an SOA. In any case, the need is to identify and integrate the required services, whilst at the same time enabling the sharing of data between services.

For their effective implementation, information management, information integration and application integration all require that the underlying information and processes be described and managed semantically, that is they are associated with a machine-processable description of their meaning. This, the fundamental idea behind the Semantic Web became prominent at the very end of the 1990s (Berners-Lee, 1999) and in a more developed form in the early 2000s (Berners-Lee *et al.*, 2001). The last half decade has seen intense activity in developing these ideas, in particular under the auspices of the World Wide Web Consortium (W3C).[1] Whilst the W3C has developed the fundamental ideas and standardised the languages to support the Semantic Web, there has also been considerable research to develop and apply the necessary technologies, for example natural language processing, knowledge discovery and ontology management. This book describes the current state of the art in these technologies.

All this work is now coming to fruition in practical applications. The initial applications are not to be found on the global Web, but rather in the world of corporate intranets. Later chapters of this book describe a number of such applications.

The book was motivated by work carried out on the SEKT project (http://www.sekt-project.com). Many of the examples, including two of the applications, are drawn from this project. However, it is not biased towards any particular approach, but offers the reader an overview of the current state of the art across the world.

1.2. THE GOAL OF THE SEMANTIC WEB

The Semantic Web and Semantic Web technologies offer us a new approach to managing information and processes, the fundamental principle of which is the creation and use of semantic metadata.

[1] See: http://www.w3.org/2001/sw/

For information, metadata can exist at two levels. On the one hand, they may describe a document, for example a web page, or part of a document, for example a paragraph. On the other hand, they may describe entities within the document, for example a person or company. In any case, the important thing is that the metadata is semantic, that is it tells us about the content of a document (e.g. its subject matter, or relationship to other documents) or about an entity within the document. This contrasts with the metadata on today's Web, encoded in HTML, which purely describes the format in which the information should be presented: using HTML, you can specify that a given string should be displayed in bold, red font but you cannot specify that the string denotes a product price, or an author's name, and so on.

There are a number of additional services which this metadata can enable (Davies *et al*., 2003).

In the first place, we can organise and find information based on meaning, not just text. Using semantics our systems can understand where words or phrases are equivalent. When searching for 'George W Bush' we may be provided with an equally valid document referring to 'The President of the U.S.A.'. Conversely they can distinguish where the same word is used with different meanings. When searching for references to 'Jaguar' in the context of the motor industry, the system can disregard references to big cats. When little can be found on the subject of a search, the system can try instead to locate information on a semantically related subject.

Using semantics we can improve the way information is presented. At its simplest, instead of a search providing a linear list of results, the results can be clustered by meaning. So that a search for 'Jaguar' can provide documents clustered according to whether they are about cars, big cats, or different subjects all together. However, we can go further than this by using semantics to merge information from all relevant documents, removing redundancy, and summarising where appropriate. Relationships between key entities in the documents can be represented, perhaps visually. Supporting all this is the ability to reason, that is to draw inferences from the existing knowledge to create new knowledge.

The use of semantic metadata is also crucial to integrating information from heterogeneous sources, whether within one organisation or across organisations. Typically, different schemas are used to describe and classify information, and different terminologies are used within the information. By creating mappings between, for example, the different schemas, it is possible to create a unified view and to achieve interoperability between the processes which use the information.

Semantic descriptions can also be applied to processes, for example represented as web services. When the function of a web service can be described semantically, then that web service can be discovered more easily. When existing web services are provided with metadata describing their function and context, then new web services can be automatically composed by the combination of these existing web services. The use of such semantic descriptions is likely to be essential to achieve large-scale implementations of an SOA.

1.3. ONTOLOGIES AND ONTOLOGY LANGUAGES

At the heart of all Semantic Web applications is the use of ontologies. A commonly agreed definition of an ontology is: 'An ontology is an explicit and formal specification of a conceptualisation of a domain of interest' (c.f. Gruber, 1993). This definition stresses two key points: that the conceptualisation is formal and hence permits reasoning by computer; and that a practical ontology is designed for some particular domain of interest. Ontologies consist of concepts (also knowns as classes), relations (properties), instances and axioms and hence a more succinct definition of an ontology is as a 4-tuple $\langle C, R, I, A \rangle$, where C is a set of concepts, R a set of relations, I a set of instances and A a set of axioms (Staab and Studer, 2004).

Early work in Europe and the US on defining ontologies languages has now converged under the aegis of the W3C, to produce a Web Ontology Language, OWL.[2]

The OWL language provides mechanisms for creating all the components of an ontology: concepts, instances, properties (or relations) and axioms. Two sorts of properties can be defined: object properties and datatype properties. Object properties relate instances to instances. Datatype properties relate instances to datatype values, for example text strings or numbers. Concepts can have super and subconcepts, thus providing a mechanism for subsumption reasoning and inheritance of properties. Finally, axioms are used to provide information about classes and properties, for example to specify the equivalence of two classes or the range of a property.

In fact, OWL comes in three species. OWL Lite offers a limited feature set, albeit adequate for many applications, but at the same time being relatively efficient computationally. OWL DL, a superset of OWL Lite, is based on a form of first order logic known as Description Logic. OWL Full, a superset of OWL DL, removes some restrictions from OWL DL but at the price of introducing problems of computational tractability. In practice much can be achieved with OWL Lite.

OWL builds on the Resource Description Framework (RDF)[3] which is essentially a data modelling language, also defined by the W3C. RDF is graph-based, but usually serialised as XML. Essentially, it consists of triples: subject, predicate, object. The subject is a resource (named by a

[2] See: http://www.w3.org/2004/OWL/
[3] See: http://www.w3.org/RDF/

URI), for example an instance, or a blank node (i.e., not identifiable outside the graph). The predicate is also a resource. The object may be a resource, blank node, or a Unicode string literal.

For a full introduction to the languages and basic technologies underlying the Semantic Web see [Antoniou and van Harmelen, 2004].

1.4. CREATING AND MANAGING ONTOLOGIES

The book is organized broadly to follow the lifecycle of an ontology, that is discussing technologies for ontology creation, management and use, and then looking in detail at some particular applications. This section and the two which follow provide an overview of the book's structure.

The construction of an ontology can be a time-consuming process, requiring the services of experts both in ontology engineering and the domain of interest. Whilst this may be acceptable in some high value applications, for widespread adoption some sort of semiautomatic approach to ontology construction will be required. Chapter 2 explains how this is possible through the use of knowledge discovery techniques.

If the generation of ontologies is time-consuming, even more is this the case for metadata extraction. Central to the vision of the Semantic Web, and indeed to that of the semantic intranet, is the ability to automatically extract metadata from large volumes of textual data, and to use this metadata to annotate the text. Chapter 3 explains how this is possible through the use of information extraction techniques based on natural language analysis.

Ontologies need to change, as knowledge changes and as usage changes. The evolution of ontologies is therefore of key importance. Chapter 4 describes two approaches, reflecting changing knowledge and changing usage. The emphasis is on evolving ontologies incrementally. For example, in a situation where new knowledge is continuously being made available, we do not wish to have to continuously recompute our ontology from scratch.

Reference has already been made to the importance of being able to reason over ontologies. Today an important research theme in machine reasoning is the ability to reason in the presence of inconsistencies. In classical logic any formula is a consequence of a contradiction, that is in the presence of a contradiction any statement can be proven true. Yet in the real world of the Semantic Web, or even the semantic intranet, inconsistencies will exist. The challenge, therefore, is to return meaningful answers to queries, despite the presence of inconsistencies. Chapter 5 describes how this is possible.

A commonly held misconception about the Semantic Web is that it depends on the creation of monolithic ontologies, requiring agreement from many parties. Nothing could be further from the truth. Of course, it is good design practice to reuse existing ontologies wherever possible, particularly where an ontology enjoys wide support. However, in many cases we need to construct mappings between ontologies describing the same domain, or alternatively merge ontologies to form their union. Both approaches rely on the identification of correspondences between the ontologies, a process known as ontology alignment, and one where (semi-)automatic techniques are needed. Chapter 6 describes techniques for ontology merging, mapping and alignment.

1.5. USING ONTOLOGIES

Chapter 7 explains two rather different roles for ontologies in knowledge management, and discusses the different sorts of ontologies: upper-level versus domain-specific; light-weight versus heavy weight. The chapter illustrates this discussion with reference to the PROTON ontology.[4]

Chapter 8 describes the state of the art in three aspects of ontology-based information access: searching and browsing; natural language generation from structured data, for example described using ontologies; and techniques for on-the-fly repurposing of data for a variety of devices. In each case the chapter discusses current approaches and their limitations, and describes how semantic web technology can offer an improved user experience. The chapter also describes a semantic search agent application which encompasses all three aspects.

The creation of ontologies, although partially automated, continues to require human intervention and a methodology for that intervention. Previous methodologies for introducing knowledge technologies into the organisation have tended to assume a centralised approach which is inconsistent with the flexible ways in which modern organisations operate. The need today is for a distributed evolution of ontologies. Typically individual users may create their own variations on a core ontology, which then needs to be kept in step to reflect the best of the changes introduced by users. Chapter 9 discusses the use of such a methodology.

Ontologies are being increasingly seen as a technology for streamlining the systems integration process, for example through the use of semantic descriptions for web services. Current web services support interoperability through common standards, but still require considerable human interaction, for example to search for web services and then to combine them in a useful way. Semantic web services, described in Chapter 10, offer the possibility of automating web service discovery, composition and invocation. This will have considerable impact in areas such as e-Commerce and Enterprise Application Integration, by

[4] http://proton.semanticweb.org/

enabling dynamic and scalable cooperation between different systems and organizations.

1.6. APPLICATIONS

There are myriad applications for Semantic Web technology, and it is only possible in one book to cover a small fraction of them. The three described in this book relate to specific business domains or industry sectors. However, the general principles which they represent are relevant across a wide range of domains and sectors.

Chapter 11 describes the key role which Semantic Web technology is playing in enhancing the concept of a Digital Library. Interoperability between digital libraries is seen as a 'Grand Challenge', and Semantic Web technology is key to achieving such interoperability. At the same time, the technology offers new ways of classifying, finding and presenting knowledge, and also the interrelationships within a corpus of knowledge. Moreover, digital libraries are one example of intelligent content management systems, and much of what is discussed in Chapter 11 is applicable generally to such systems.

Chapter 12 looks at an application domain within a particular sector, the legal sector. Specifically, it describes how Semantic Web technology can be used to provide a decision support system for judges. The system provides the user with responses to natural language questions, at the same time as backing up these responses with reference to the appropriate statutes. Whilst apparently very specific, this can be extended to decision support in general. In particular, a key challenge is combining everyday knowledge, based on professional experience, with formal legal knowledge contained in statute databases. The development of the question and answer database, and of the professional knowledge ontology to describe it, provide interesting examples of the state of the art in knowledge elicitation and ontology development.

The final application, in Chapter 13, builds on the semantic web services technology in Chapter 10, to describe how this technology can be used to create an SOA. The approach makes use of the Web Services Modelling Ontology (WSMO)[5] and permits a move away from point to point integration which is costly and inflexible if carried out on a large scale. This is particularly necessary in the telecommunications industry, where operational support costs are high and customer satisfaction is a key differentiator. Indeed, the approach is valuable wherever IT systems need to be created and reconfigured rapidly to support new and rapidly changing customer services.

[5] See http://www.wsmo.org/

1.7. DEVELOPING THE SEMANTIC WEB

This book aims to provide the reader with an overview of the current state of the art in Semantic Web technologies, and their application. It is hoped that, armed with this understanding, readers will feel inspired to further develop semantic web technologies and to use semantic web applications, and indeed to create their own in their industry sectors and application domains. In this way they can achieve real benefit for their businesses and for their customers, and also participate in the development of the next stage of the Web.

REFERENCES

Antoniou G, van Harmelen F. 2004. *A Semantic Web Primer*. The MIT Press: Cambridge, Massachusetts.

Berners-Lee T. 1999. *Weaving the Web*. Orion Business Books.

Berners-Lee T, Hendler J, Lassila O. 2001. The semantic web. In *Scientific American*, May 2001.

Charlesworth I. 2005. Integration fundamentals, Ovum.

Davies J, Fensel D, van Harmelen F (eds). 2003. *Towards the Semantic Web: Ontology-Driven Knowledge Management*. John Wiley & Sons, Ltd. ISBN: 0470848677.

Drucker P. 1999. Knowledge worker productivity: the biggest challenge. *California Management Review* 41(2):79–94.

Fensel D, Hendler JA, Lieberman H, Wahlster W (eds). 2003. *Spinning the Semantic Web: Bringing the World Wide Web to its Full Potential*. MIT Press: Cambridge, MA. ISBN 0-262-06232-1.

Gruber T. 1993. A translation approach to portable ontologies. *Knowledge Acquisition* 5(2):199–220, http://ksl-web.stanford.edu/KSL_Abstracts/KSL-92-71.html

Holsapple CW Eds. 2002. *Handbook on Knowledge Management*. Springer: ISBN:3540435271.

Lyman P, *et al*. 2005. How Much Information? 2003, School of Information Management and Systems, University of California at Berkeley, http://www.sims.berkeley.edu/research/projects/how-much-info-2003/

Morello D. 2005. The human impact of business IT: How to Avoid Diminishing Returns.

Staab S, Studer R (Eds). 2004. *Handbook on Ontologies. International Handbooks on Information Systems*. Springer: ISBN 3-540-40834-7.