The Knowledge Engineering Review, Vol. 20:2, 191–196.
doi:10.1017/S0269888905000500 Printed in the United Kingdom

Towards Semantic Web agents: Knowledge Web and AgentLink

TERRY PAYNE[1] and VALENTINA TAMMA[2]

[1]*School of Electronics and Computer Science, The University of Southampton, Highfield, Southampton SO17 1BJ, UK; e-mail: trp@ecs.soton.ac.uk*
[2]*Department of Computer Science, University of Liverpool, Liverpool L69 7ZF, UK; e-mail: valli@csc.liv.ac.uk*

Abstract

This paper presents an overview on the role of agents in the Semantic Web, which was the topic of the AgentLink Technical Forum Group on 'Semantic Web Agents', aimed at fostering closer collaboration between the European communities working in these areas. The paper is structured in three main sections. In the first, we argue how agents are an essential component of the Semantic Web, then we provide a brief history, by no means comprehensive, of how the Semantic Web vision—that includes agents—has evolved over the past fifteen years. We then conclude, reporting on the topics presented and discussed during the Technical Forum.

1 What is the Semantic Web

Semantic Web technology (Fensel *et al*., 2003) is becoming increasingly popular because it promises to add value to the current Web without requiring any fundamental changes to the Web infrastructure that is currently in place. In its *manifesto* by Berners-Lee, Hendler, and Lassila (Berners-Lee *et al*., 2001), the authors give the following definition: 'The Semantic Web is not a separate Web but an extension of the current one, in which information is given well-defined meaning, better enabling computers and people to work in cooperation'. Therefore, the Semantic Web extends the current Web so that it is more easily accessible by computer programs, and especially *software agents*. Indeed, the Web of today is intended primarily for human consumption, and is based on visual markup languages (such as HTML) intended for human users. In contrast, the Semantic Web (SW) is intended to be utilized and understood by software agents, and is based on machine-processable languages such as RDF(S), DAML+OIL (DAML, 2001), and OWL (McGuinness & van Harmelen, 2003), which permit knowledge to be embedded within Web pages. Software agents are able to process the knowledge expressed in these semantic markup languages, and can thus offer services which make use of or retrieve this knowledge.

The Semantic Web is based on the harmonious combination of *ontologies*, *reasoning mechanisms*, and *agents*. Ontologies are an explicit, formal specification of a shared conceptualization of a domain (Studer *et al*., 1998); they provide a machine readable, and agreed upon representation of an abstraction of some phenomenon. Knowledge represented in ontologies can be enriched through reasoning mechanisms. Ontologies written in OWL can take advantage of DL-based reasoners, such as Racer (Haarslev & Moller, 2001) or Pellet,[1] which permit the classification of new instances or inference of new facts. In some cases, ontologies can be complemented by *axioms*, statements that are always true and that are used to constrain the meaning of concept definitions in the

[1] Pellet OWL Reasoner; `http://www.mindswap.org/2003/pellet/index.shtml`.

ontologies. These axioms can be represented through *rules* built on top of ontological definitions, and permit enhanced representation and reasoning capabilities. Rules play an integral role in the Semantic Web, and the rule layer is meant to provide an additional mechanism to support the deduction of new knowledge and combination of information. The effort to define rule languages, and to enable interoperability between major commercial rule systems has been undertaken by the SWRL Initiative (Horrocks *et al*., 2004), preceded by the RuleML initiative (Grosof *et al*., 2003b). Some approaches, e.g. SweetJess (Grosof *et al*., 2003a) and SweetProlog (Laera *et al*., 2004), allow agents committing to an OWL ontology to reason over its instances using SWRL or RuleML rules translated into some Logic Programming Language, such as Prolog or Jess, that offer efficient automatic reasoning. Agents can thus use the domain knowledge represented in ontologies, together with a set of rules, in order to accomplish some tasks on behalf of users.

2 A brief history of agents and the Semantic Web

The use of ontologies (with axioms) within multi-agent systems is a topic that has recently received much attention. On the one hand, the agent paradigm is successfully employed in those applications where autonomous, loosely coupled, heterogeneous, and distributed systems need to interoperate in order to achieve a common goal. On the other hand, ontologies have established themselves as a powerful tool to enable knowledge sharing, and a growing number of applications have benefited from the use of ontologies as a means to achieve semantic interoperability among heterogeneous, distributed systems. In principle, ontologies and agents are a 'match made in heaven', which was carefully planned but has so far failed to happen. In the early 1990s the Knowledge Sharing Initiative (Neches *et al*., 1991) aimed at achieving interoperability between different applications, where interoperability was to be achieved at different levels (syntactic, semantic, etc.). This research effort introduced the notion of ontology as a means of achieving *semantic* interoperability between different applications committing to different knowledge bases. These knowledge bases could differ in the language used to represent knowledge, but also in the way they conceptualized the domain.

The same research effort also concentrated on the notion of interoperability between different external applications, by defining KQML (Knowledge Query Manipulation Language), a language that specifies the interface between agents providing knowledge services, and external applications requiring these services. KQML prescribes the way an agent handles certain fragments of knowledge by means of speech act performatives (Searle, 1983), such as REQUEST and INFORM. The language specifies the syntax of the messages to exchange, but not the semantics. The semantics are specified by ontologies; agents commit to the same shared ontology, which provided them with a understanding of this knowledge (the object in the messages represented by the performatives). Domain interoperability is made possible through libraries of shared domain ontologies. Figure 2 illustrates the conceptual architecture enabling interoperability between different applications.

The type of approach proposed by the Knowledge Sharing Initiative was accepted in principle, but also has a few flaws. The approach was perceived as cumbersome, for instance, the language chosen to represent knowledge, KIF, was based on first-order logic, hence ontology modelling was restricted to knowledge practitioners and there was limited tool support. In addition, KIF is not decidable, thus implying that agents who need to reason with KIF representation might never terminate their tasks.

Despite its limitations, the conceptual architecture proposed by the Knowledge Sharing Initiative has become the recommended approach in the agent community (Huhns & Singh, 1997). FIPA (Foundation for Intelligent Physical Agents) is a non-profit organization aimed at producing standards for the interoperation of heterogeneous software agents. FIPA introduced a standard agent communication language, FIPA ACL (FIPA, 1997) that FIPA compliant agents use in order to communicate with other agents. FIPA-ACL resembles KQML, in that it is also based on speech act theory, and it only defines the syntax of the performatives while the semantics is left unspecified and can be defined in terms of some shared ontology to which the agents commit. During the 1990s agent architectures were applied in open environments, where agents are free to join and leave

Figure 1 The Semantic Web

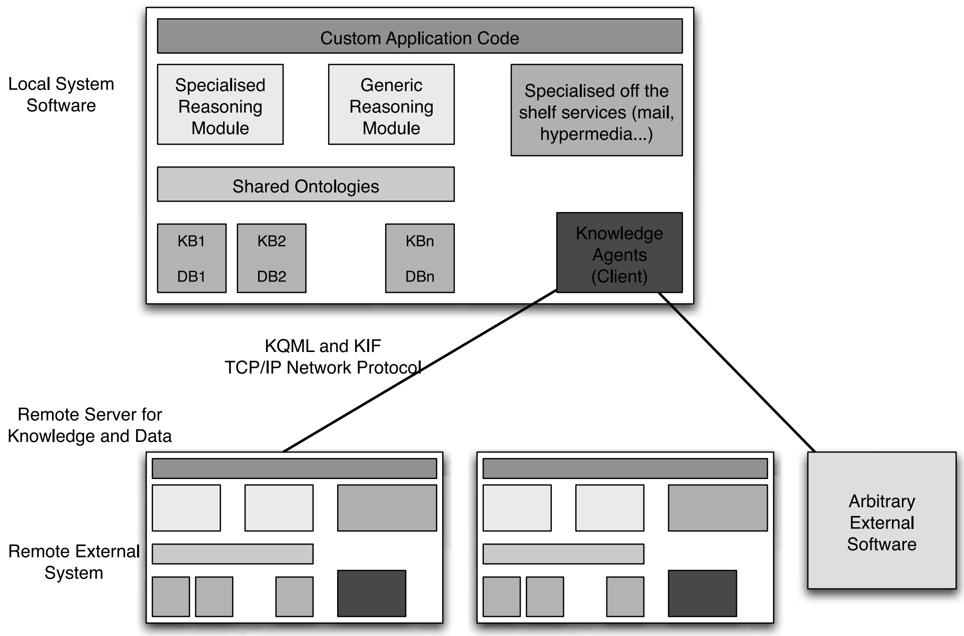

Figure 2 The conceptual architecture of the knowledge sharing initiative

interoperation at any time. In open environments coordination is achieved through middle agents (Klusch & Sycara, 2001), whose role is to provide mediation services for the processing of agent capabilities and service descriptions, semantic interoperation between agents and external systems, and management of data and knowledge. Many of the mediators proposed in the literature are based on common and shared representation of agent capabilities, often represented in ontologies, such as in LARKS (Sycara *et al*., 1998). A comprehensive review of middleware services is provided in Klusch & Sycara (2001). This view has also been taken in programmes funded by the US Defence Advanced Research Projects Agency (DARPA), such as the CoABS (Control of Agent Based Systems) project, aiming at providing an infrastructure for permitting interoperability in large-scale, cooperative teams of interacting agents, each offering different capabilities. Thus, this infrastructure permits the integration at run time of heterogeneous systems that have been separately developed. The discovery of these systems is enabled by taxonomic descriptions of common functionalities. Where gaps might exist between the agents, translation services are used to provide greater interoperability by seamlessly filling in the pieces (Kahn & Cicalese, 2002). More recently the CoABS infrastructure has been proposed as a possible means to provide agent-based support to Internet-based grids management (Overeinder *et al*., 2002).

At the same time DARPA was also promoting research on large-scale knowledge representation, through the PARKA project. PARKA is a frame-based knowledge representation system which uses the massive parallelism of the Connection Machine (CM) to perform certain inferences extremely rapidly. PARKA performs property inheritance, set-theoretic inferences, and inferences on partonomic relations. The representation framework proposed by PARKA was successfully employed in mobile robot applications (Seeliger & Hendler, 1995). One of the outcomes of the PARKA project was a language for representing knowledge on the Web, SHOE (Heflin & Hendler, 2000), one of the first efforts to define a markup language for ontologies.

From efforts such as CoABS and PARKA, the need for a language to represent information that is machine readable and understandable emerged. The reason is that the World Wide Web contains a large amount of information which is expanding at a rapid rate. However, this information is represented in a human friendly way, but is not easily accessible to machines. The DAML project (DARPA Agent Markup Language) was launched in 2000 and aimed at developing a markup language to extend XML, which could be use to define information objects and to capture the relationships between objects by means of schemas or ontologies (DAML, 2001). At the same time, in Europe, several European projects were addressing a similar problem, that of developing a language for adding semantic descriptions to the WWW. In particular, Ontobroker developed OIL (Ontology Inference Layer), a language for a Web-based representation and inference layer for ontologies, which combines the widely used modelling primitives from frame-based languages with the formal semantics and reasoning services provided by description logics (Fensel *et al*., 2000). These two efforts lately merged in the DAML+OIL framework that then became OWL, the Ontology Web Language (OWL). In the OWL framework, concepts can be defined within extensible open ontologies published in OWL (or DAML+OIL), with instances and properties being asserted at arbitrary locations across the Internet. Software components, and in particular multi-agent systems, provide methodologies for consuming, processing, creating and disseminating both information and knowledge within a distributed framework. Depending on the type of applications in which they are applied, the ontological representations and the agents reasoning capabilities range from very simple to very complex.

Currently, research on ontologies in multi-agent systems covers a wide range of topics but can mainly be grouped into three categories: modelling principles for building and reasoning with ontologies for agents, semantic interoperability between different agents, and applications of ontologies in agent systems. An example of the kind of research carried out can be found in Tamma *et al*. (2005).

3 Where we are: the Technical Forum

The Semantic Web thus presents a different perspective in the way distributed knowledge can be viewed, in terms of creation, consumption, dissemination, and management. Concepts can be defined within extensible, open ontologies published using standard protocols with instances and properties being asserted at arbitrary locations across the Internet. Multi-agent systems, in contrast, provide methodologies for consuming, processing, creating and disseminating both information and knowledge within a distributed framework. Moreover, the construction and co-ordination of heterogeneous agent communities requires complex management of data, the interpretation of which may be ambiguous without semantic structure or reasoning. The aim of the SWA Technical Forum Group—'Towards Semantic Web Agents: Knowledge Web and AgentLink', was to explore the synergy between the two research communities, and to examine how the two largest EU projects supporting research in this area, KnowledgeWeb[2] and AgentLink,[3] could exploit this. Three half-day sessions were planned as part of the Technical Forummeeting, to address three questions.

[2] http://knowledgeweb.semanticweb.org/.
[3] http://www.agentlink.org.

- How can KnowledgeWeb and AgentLink benefit from working together?
- What can the Semantic Web do for Agents?
- What can Agent Technology do for the Semantic Web?

During the first session, talks were given from senior members of the management committees of both projects. The aims and structures of each project were presented (by Peter McBurney and Jerome Euzenat), and it was soon apparent that both projects shared similar aims and challenges, in terms of education, industrial outreach, and mapping out the current and future key issues of each project. Education was viewed as important by both communities, and Joerg Diederich presented the VISWE archive—Virtual Institute for Semantic Web Education—which is a joint project by both KnowledgeWeb and the REWERSE Network of Excellence[4] to support education across the EU. The remaining research-led talks demonstrated how each community could benefit from the other. If any two heterogeneous systems that ground their knowledge in different ontologies are to interoperate, an alignment is first necessary between the two ontologies before subsequent transactions can take place. Jerome Euzenat argued that an agreement needs to be established at the level of the ACL performatives used to communicate (i.e. the envelope of the message) but also at the content level, by defining the semantics of the body of the message. The need to find an alignment between heterogeneous ontology that may lead to defining mappings either statically or dynamically was also considered in the presentations by Aleksander Pivk and Nuno Silva. The risk of this complex process being trivialized or dismissed is great, especially when there is a lack of understanding of the issues involved. The Semantic Web Space project, based on a triple-variant of Linda spaces, proposes a distributed space where agents are able to share triples by means of an RDF-based extension of Linda. The invited talk focused on the relation between agents and (Semantic)Web services. This talk characterized agents versus Web services, Semantic Web services and Grid services, arguing how the service community can benefit from some of the more typical agent characteristics, such as the ability to act towards the achievement of a goal, and hence evaluate the utility of a specific course of action. Related to the invited talk was the contribution by Monika Solanki, who presented some initial issues on composition of agent-based services on the Semantic Web. The other talks presented during the second day were all research oriented, and concentrated on the use of agents for semantic brokerage (Julian Padget), and the need to combine agents, Semantic Web technology, and natural language techniques for the annotation of Web resources (Patrick Storms).

The second day also hosted the panel session focussing on two chief topics: 'What agents can do for the Semantic Web' and 'What the Semantic Web can do for agents'. The panellists agreed on the necessity to isolate commonalities between the two areas and to separate what is unique and can be contributed to the other field. Among the unique characteristics of agents, Jerome Euzenat identified the well-defined notions of Agent Communication Languages and interaction protocols. Terry Payne advocated the use of agents to provide smart, service-based support for autonomous Semantic Web tools, and well-automated discovery mechanisms for advertising and locating resources within an open framework, established trust and reputation frameworks, and proactive support for fact maintenance. Jerome Euzenat argued that the unique contribution of the Semantic Web is the notion of ontologies. In particular, Jerome stressed the importance of clearly identifying the type of problems tackled by agents, and that the Semantic Web can help with heterogeneity and communication problems. Terry Payne argued that multi-agent systems benefit from Semantic Web technologies that provide semantic interoperability between heterogeneous, independent sources, which permits communication by helping resolving differences; the provision of a shared, distributed instance space; and reasoning support for conceptual negotiation and argumentation. Monika Solanki argued that the discussion cannot ignore the Web service perspective, and in this perspective Semantic Web technology represents the bridge that creates the synergy between agents and Web services: the Semantic Web can enable the realization of Web services by providing

[4] `http://rewerse.net/`.

semantic content and syntactic unification among the several protocols and languages used to model Web services. Agents have the potential to provide autonomous and stateful Web services.

The industrial viewpoint was presented by Patrick Storms. In his opinion, the main contribution from the agent community is the automation of tedious and labour intensive tasks, such as the annotation of Web documents, or the provision of housekeeping services for the Semantic Web. Semantic Web technology can also support agents when dealing with unstructured data, by providing means for semantic reconciliation. In addition, the Semantic Web can become an enabler for new agent-based services.

Acknowledgements

We would like to express our appreciation to the attendees of the AgentLink III Technical Forum on Semantic Web Agents in Ljubljana, 2005, and to both AgentLink and KnowledgeWeb for their support.

References

Berners-Lee, T *et al*., 2001, The Semantic Web. *Scientific American* (284), 35–43.

DAML, 2001, The DARPA Agent Markup Language, http://www.daml.org/.

Fensel, D *et al*., 2000, OIL in a nutshell. In RD *et al*. (eds.), *Knowledge Acquisition, Modeling, and Management, Proceedings of the European Knowledge Acquisition Conference (EKAW-2000) (Lecture Notes in Artificial Intelligence)*. Berlin: Springer.

Fensel, D *et al*. (eds.), 2003, *Spinning the Semantic Web: Bringing the World Wide Web to its Full Potential*. Cambridge, MA: The MIT Press.

FIPA, 1997, FIPA ACL 1997: Specification Part 2—Agent Communication Language. Technical report.

Grosof, B *et al*., 2003a, SweetJess: inferencing in situated courteous RuleML via translation to and from Jess rules, `http://ebiz.mit.edu/bgrosof/paps/sweetjess-wp-050203.pdf`.

Grosof, B *et al*., 2003b, Description logic programs: combining logic programs with description logic. In *Proceedings of the Twelfth International World Wide Web Conference (WWW 2003)*. New York: ACM Press, pp. 48–57.

Haarslev, V and Moller, R, 2001, RACER system description. In *Proceedings of the International Joint Conference on Automated Reasoning, IJCAR'2001*.

Heflin, J and Hendler, J, 2000, Dynamic ontologies on the Web. In *Proceedings of the Seventeenth National Conference on Artificial Intelligence (AAAI-2000)*. AAAI/MIT Press, pp. 443–449.

Horrocks, I *et al*., 2004, SWRL, `http://www.w3.org/Submission/2004/SUBM-SWRL-20040521/`.

Huhns, M and Singh, M, 1997, Ontologies for agents. *IEEE Internet Computing, Special Issue on Agents on the Web* November–December.

Kahn, M and Cicalese, CDT, 2002, The CoABS grid. In *Proceedings of the Goddard–JPL Workshop on Radical Agent Concepts*, pp. 16–18.

Klusch, M and Sycara, K, 2001, Brokering and matchmaking. In Omicini, A, Zambonelli, F, Klusch, M and Tolksdorf, R (eds.), *Coordination of Internet Agents*. Berlin: Springer.

Laera, L *et al*., 2004, Sweetprolog: a system to integrate ontologies and rules. In *Proceedings of ISWC 2004 workshop on Rules and Rule Markup Languages for the Semantic Web*.

McGuinness, DL and van Harmelen, F, 2003, OWL, Web Ontology Language, http://www.w3.org/TR/2003/CR-owlfeatures- 20030818/.

Neches, R *et al*., 1991, Enabling technology for knowledge sharing. *AI Magazine* **12**(3), 36–56.

Overeinder, B *et al*., 2002, Multi-agent support for Internet-scale Grid management. In Rana, O and Schroeder, M (eds.), *Proceedings of the AISB'02 Symposium on AI and Grid Computing*, pp. 18–22.

Searle, J, 1983, *Intentionality*. Cambridge: Cambridge University Press.

Seeliger, O and Hendler, J, 1995, Supervenient hierarchies of behaviors: from robot vacuuming to space telerobotics. *Journal for Experimental and Theoretical Artificial Intelligence*.

Studer, R *et al*., 1998, Knowledge engineering, principles and methods. *Data and Knowledge Engineering* **25**(1–2), 161–197.

Sycara, K *et al*., 1998, Dynamic service matchmaking among agents in open information systems. *ACM SIGMOD Record. Special Issue on Semantic Interoperability in Global Information Systems*.