Computers in Industry 61 (2010) 729–741

Contents lists available at ScienceDirect

Computers in Industry

journal homepage: www.elsevier.com/locate/compind

Semantic Web computing in industry

John G. Breslin *, David O'Sullivan, Alexandre Passant, Laurentiu Vasiliu

School of Engineering and Informatics, National University of Ireland, Galway, Ireland

ARTICLE INFO

Article history:
Available online 22 June 2010

Keywords:
Semantic Web
Semantic Web Services
Knowledge Management
E-Commerce
Enterprise Application Integration

ABSTRACT

The Semantic Web has attracted significant attention during the last decade. On the one hand, many research groups have changed their focus towards Semantic Web research and research funding agencies particularly in Europe have explicitly mentioned Semantic Web in their calls for proposals. On the other hand, industry has also begun to watch developments with interest and a number of large companies have started to experiment with Semantic Web technologies to ascertain if these new technologies can be leveraged to add more value for their customers or internally within the company, while there are already several offers of vendors of Semantic Web solutions on the market. The essence of the Semantic Web is to structure Web-based information to make it more interoperable, machine-readable and thereafter to provide a means to relate various information concepts more easily and in a reusable way. The Semantic Web acts as an additional layer on the top of the Web, and is built around explicit representations of information concepts and their relationships such as ontologies and taxonomies. Furthermore, Semantic Web technologies are not only valuable on an open environment like the Web, but also in closed systems such as in industrial settings. Hence, these technologies can be efficiently deployed for domains including Web Services, Enterprise Application Integration, Knowledge Management and E-Commerce, fulfilling existing gaps in current applications. This paper focuses on this synthesis between Semantic Web technologies and systems problems within industrial applications. There will be a short review of Semantic Web standards, languages and technologies followed by a more detailed review of applications of Semantic Web computing in industry. The paper covers theoretical considerations as well as use cases and experience reports on the topic, and we also present some current challenges and opportunities in the domain.

1. Introduction

The Semantic Web is now becoming a well-established branch of computer science and software engineering with its own standards, languages, technologies and applications. It is also a foundation for what is termed 'Web Science', where the Web itself is the object of a dedicated science of its own when it is deployed in a wide range of domains.[1] There are a number of research institutes now feeding new knowledge into the associated research community, and a large number of new and existing industries are deploying Semantic Web techniques to provide goods and services to customers. While the World Wide Web and its associated technologies and applications have become a 'disruptive technology' over a relatively short period of time, it remains to be seen whether the Semantic Web with its related new technologies and applications will do the same. There are nevertheless some encouraging indications. The number of new business start-ups that now deploy Semantic Web technologies has become noticeable. Web 2.0 companies such as Freebase, Faviki and Zemanta have embraced Semantic Web technologies. The New York Times also identified commercial industries around the world that are using Semantic Web technologies as part of their core business offerings to customers [1]. Giants such as Oracle, Vodafone, Amazon, Adobe, Microsoft, Yahoo and Google are now experimenting with Semantic Web technologies to provide new value to customers [2], with some recent efforts including the Yahoo! SearchMonkey search engine [3] and Google's support in indexing structured RDF information from the Web.[2]

The core of the Semantic Web contains a number of fundamental formal models, languages and technologies for interoperability and reuse of information, including RDF, RDFS, the OWL family of languages, the WSML family of languages and SPARQL. Semantic Web Services build on the Semantic Web and previous work regarding Web Services to power semi-automated or automated interoperable applications. In this paper, we will

* Corresponding author.
E-mail address: John.Breslin@nuigalway.ie (J.G. Breslin).

[1] http://webscience.org.

[2] http://googlewebmastercentral.blogspot.com/2009/05/introducing-rich-snippets.html.

0166-3615/$ – see front matter © 2010 Elsevier B.V. All rights reserved.
doi:10.1016/j.compind.2010.05.002

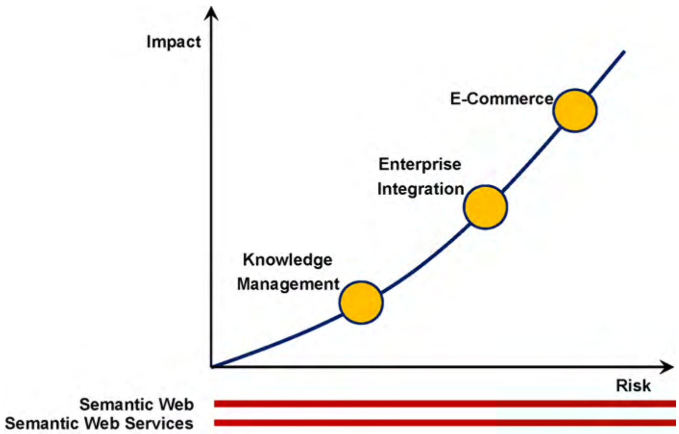

Fig. 1. Industrial applications for the Semantic Web and Semantic Web Services.

describe the Semantic Web and Semantic Web Services technologies which act as foundational layers for a variety of semantic industrial applications. We will detail three key areas for applications: Knowledge Management, Enterprise Application Integration and E-Commerce. These three application areas are shown in Fig. 1, where the *Y*-axis represents the reward or potential impact for semantics within industry, and the *X*-axis represents the risk involved in implementing change to existing technologies. For example, deploying semantics in the field of Knowledge Management may not yield the same rewards as applications in e-commerce, but the risks involved in successful e-commerce deployment are much higher. Each of these three major application domains is built upon two major areas of research—the Semantic Web and Semantic Web Services. We begin our review by looking at these two areas and then looking at each of the major application domains in turn.

In the following section, we will introduce the main directions of development within the Semantic Web along with their main technologies, tools and achievements to date, as well as describing various limitations and possible future developments.

2. The Semantic Web

The 'Semantic Web' can be thought of as the next generation of the Web where computers that can aid humans with their daily web-related tasks as more meaningful structured information is added to the Web (manually and automatically) [4]. For example, using a combination of facts like "John works_at NUI Galway", "Mary knows John", "a Person works_at an Organisation", and "a Person knows a Person", you can allow computers to answer relatively straightforward questions like "Find me all the people who know others who work at NUI Galway" which at this moment is quite difficult to do without significant manual processing of the information returned from search results. The Semantic Web represents these facts through the use of metadata that is associated with Web resources, and behind this metadata there are specific vocabularies or 'ontologies' [5] that describe what are the semantics (or meaning) of this metadata and how it is all related to each other.

Metadata can be thought of as 'data about data'. Similar to how librarians traditionally put information about books into catalogues or library cards, metadata on the Web commonly refers to descriptive information about Web resources that can support a wide range of operations [6] ranging from retrieving to re-contextualising content. Metadata elements are used to give structure to the description of a resource. For example in an educational course, metadata elements will include title, description, keywords, author, educational level, version, location, language, date created, and so on. RDF (Resource Description Framework) is used to express metadata about resources [7] while these resources are defined using URIs (Uniform Resource Identifier) such that they are provided with unique and non-ambiguous identifiers at Web-scale, enabling interoperability between various applications. Led by the W3C consortium, RDF is supported by a wide range of stakeholders ranging from digital librarians to B2B industries and has achieved significant industrial momentum.

RDF consists of two aspects: a graph-based abstract model for the data (made up of multiple statements, or triples) and the RDF syntax (with a variety of serialisations to represent these triples in a computer-readable form such as N3, Turtle, RDF/XML or RDFa which allows RDF annotations to be directly embedded within XHTML pages). For example, to say that Alice knows Bob, we could use the Notation3 (N3) syntax for the corresponding RDF triple: "http://example.org#Alice http://xmlns.com/foaf/0.1/knows http://example.org/#Bob". All triples are in the form of a directed graph, from subject via a directed arc (the predicate) to an object. In the previous example, Alice would be the subject, the 'knows' relationship is the predicate and Bob is the object. URIs are normally used to give identifiers to the subject, predicate and object, but the object may sometimes be a literal or text string if an attribute is to be assigned to a subject, e.g. "http://example.org#Alice http://xmlns.com/foaf/0.1/name 'Alice Cooper'". A sample RDF graph model is shown in Fig. 2.

Further structure is provided by a metadata schema or ontology (e.g. as shown in the bottom layer of Fig. 2). For example, if there is metadata about a soccer team, an underlying ontology will say that a soccer team always has a goalkeeper and always has one and only one manager, so each metadata entry for a soccer team should have that information. Ontologies are formal and consensual specifications of conceptualisations that provide a shared and common understanding of a domain [5]. In order to deploy ontologies on the Web, two languages have been put forward as standard proposals by the W3C, namely RDFS and OWL. RDF schema (RDFS) is commonly used for the definition of RDF ontologies (and written in RDF) on the Semantic Web [8]. Some of the more popular Semantic Web lightweight ontologies include FOAF (Friend-of-a-Friend, for social networks) [9], Dublin Core (for resources online or in libraries), SIOC (Semantically-Interlinked Online Communities, for online communities and content) [8], and the Geo vocabulary[3] (for geographic locations). Recently, Bizer et al. [10] provided a list of popular and core vocabularies that people should use when publishing data on the Semantic Web as well as some best practices for publishing RDF data on the Web.

While popular, RDFS is somewhat limited in various regards. In order to overcome some of the limits of RDFS, ontology developers can use OWL (the Web Ontology Language) [11] (currently being revised towards OWL2)[4] to define more precise axioms within their ontologies, for example, transitivity of some properties (e.g. in an "ancestor" property), symmetry (e.g. "sibling") or cardinality constraints (such as the "has one and only one manager" in the previous example). In addition, ontologies also act as a support for reasoning systems, both to derive new facts or to check the consistency of the model. OWL provides three increasingly expressive sublanguages designed for use by specific communities of implementers and users: OWL Lite, OWL DL and OWL Full. OWL Lite supports those users primarily needing a classification hierarchy with simple constraints. OWL DL supports those users who want the maximum expressiveness while retaining computational completeness, while OWL Full is meant for users who want

[3] http://www.w3.org/2003/01/geo/.

[4] http://www.w3.org/2007/OWL/wiki/OWL_Working_Group.

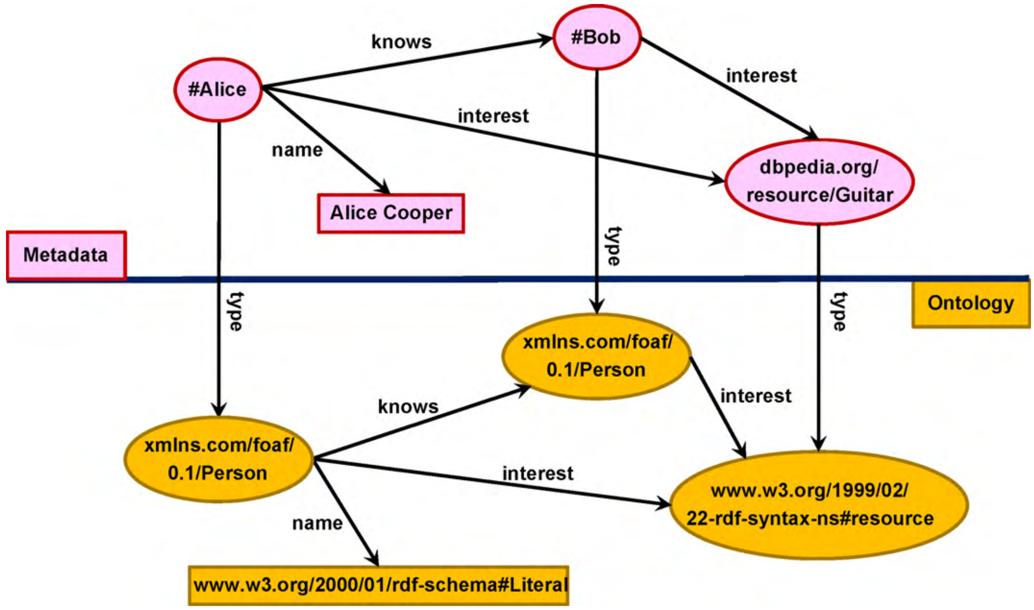

Fig. 2. Metadata and ontologies.

maximum expressiveness and the syntactic freedom of RDF with no computational guarantees.

Once this metadata has been published, using RDF(S)/OWL, query languages are required to make full use of it. SPARQL (Protocol And RDF Query Language) aims to satisfy this goal and provides both a query language and a protocol for accessing RDF data [12]. SPARQL can be thought of as the SQL of the Semantic Web, and offers a powerful means to query RDF triples and graphs. As Tim Berners-Lee stated[5]: "Trying to use the Semantic Web without SPARQL is like trying to use a relational database without SQL. SPARQL makes it possible to query information from databases and other diverse sources in the wild, across the Web." SPARQL is a graph-querying language, which means that the approach is different than SQL where people deal with tables and rows. Since query patterns are based on the RDF graph model, advanced queries are made possible such as "find every person who knows someone who knows someone else interested in Semantic Web technologies". In addition, to overcome some of the current limitations of the language, the SPARQL Working Group in W3C is now working on updating the language and considering some new features such as aggregates and updates [13].

When data is represented using RDF and can be accessed with SPARQL, queries can be created that are relevant to a particular organisation, e.g. "show me the most popular or least popular reports", or "show me any reports that used some of my data". According to Eric Miller from Zepheira [14], this can bring organisations into a "Linked Enterprise Data" (LED) framework,[6] a parallel idea to the Linking Open Data[7] initiative (a community project focusing on providing interlinked RDF data from existing open sources, leading to the availability of billion of resources and triples on the Web, and based on the Linked Data principles defined in [15]). LED aims to both expose and link enterprise data, while showing that there are benefits in terms of solutions that can be made available immediately.

Other components are required to achieve the complete Semantic Web vision, including proof, trust and user interfaces and applications, leading to the Semantic Web stack designed by the W3C as depicted in Fig. 3. Some of these components are described under the discussion on Semantic Web Services.

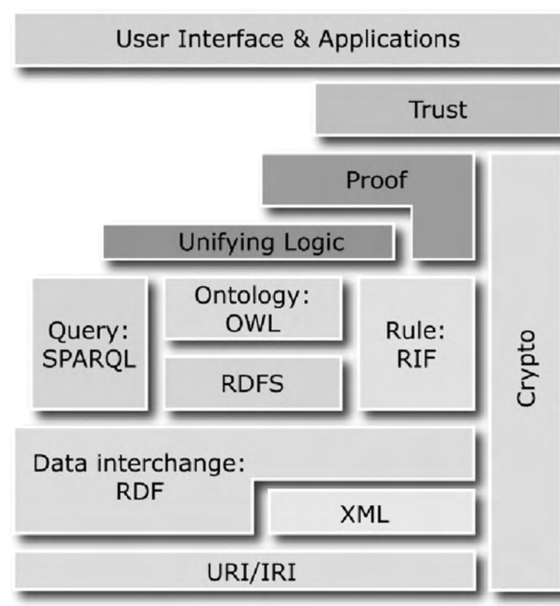

Fig. 3. The Semantic Web stack.[27]

3. Semantic Web Services

The Semantic Web Services (SWS) domain appeared in 2004 with the idea of combining the Semantic Web with Web Services technologies in order to increase automation and business execution in a digital environment. A strong motivation was to address the huge integration problems that IT industry faces continuously, including data interoperability issues, formats heterogeneity and many other issues that appear when several IT systems need to be interconnected. Their combination was a natural way of enhancing the power of possible semi-automated or

[5] http://www.w3.org/2007/12/sparql-pressrelease.
[6] A summary of the presentation is given at http://tinyurl.com/ericmiller.
[7] http://linkeddata.org/.
[27] http://www.w3.org/2007/03/layerCake.png.

automated applications in order to save time and operational costs. Semantic technology is used to add 'meaning' (machine processable declarative features) to the specifications and implementations of Web Services, to make possible the integration of distributed autonomous systems, while having independently designed data and behaviour models. Data definitions, behaviour and system components are designed in a machine-understandable form by using ontologies defined in new languages (e.g. WSMO—Web Services Modelling Ontology), providing the basis for reducing the need for human intervention in system integration processes. Web services (WS) are used in the context of automatically publishing functions or content to the rest of the world on the Web. They are application components and they communicate by using open protocols. The basic WS platform is made of XML and HTTP. Within their original design, they were aimed to be self-contained and self-describing. UDDI (Universal Description, Discovery and Integration) repositories together with SOAP (Simple Object Access Protocol) and WSDL (Web Services Description Language) make up the constituent elements of WS. More concretely, a WS presents (i) a capability that is a functional description of a Web Service, describing constraints on the input and output of a service through the notions of pre-conditions, assumptions, post-conditions, and effects and (ii) interfaces that specify how the service behaves in order to achieve its functionality. A service interface consists of a choreography that describes the interface for the client–service interaction required for service consumption, and an orchestration that describes how the functionality of a Web Service is achieved by aggregating other Web Services.

Some key languages can be used for describing Semantic Web Services including OWL [11], OWL-S [16] and WSMO/WSML (Web Services Modelling Ontology/Language) [17]. The W3C's Semantic Web stack in Fig. 3 shows a layering of technologies that can be traced bottom-up from XML, RDF, SPARQL, OWL, a logic layer and some final layers on top. While an evolutionary path can be followed from XML towards RDF and OWL,[8] the WSML language has been created independently of the previous languages (but it sits at the same level as OWL) with the goal of overcoming some limitations of previous technologies in relation to Web Services. Description logic [18] is being used in the OWL family of languages while description logic, first-order logic [19] and logic programming [20] are all used in the WSML family of languages.

3.1. OWL/OWL-S

OWL was one of the first languages to be considered for use in combination with Web Services. OWL builds on XML's ability to define customised tagging schemes and on RDF's flexible approach to representing data. If machines are expected to perform useful reasoning tasks on these documents, the language must sometimes go beyond the basic semantics of RDF schema. OWL was designed to meet this need for a full Web Ontology Language, and is part of a growing stack of W3C recommendations related to the Semantic Web.

OWL-S was created with the purpose of being a semantic markup language for Web Services and was submitted to the W3C in November 2004 [21]. The motivation behind this was to ensure automatic Web Service discovery, invocation, composition and interoperation. In order to provide such functionality for Web Services consumption, it was necessary to be able to declare 'what does a service provide', 'how a service can be used' and 'how does the interaction process occur'. For these functionalities, concepts were created such as 'service profile', 'service model' and 'service grounding'. Within the W3C specification on OWL-S, details can be found for each concept regarding their features and outlining how the OWL language can be used to address the profile description of a service, how to model a process to be executed and how to ground a service—protocols, serialisation, transports and addressing.

OWL and OWL-S have become quite popular in the academic research around the world. The missing part from this approach was a standardised execution environment that can process and handle the transactions of OWL ontologies and web service descriptions so as to realise discovery, invocation and composition. Interested users needed to build execution environments for themselves in order to execute the applications they were interested in. The industry pick-up has been quite modest and to date there are no major usages or killer applications of OWL/OWL-S technologies. This may be due to the high complexity of the technology itself, the lack of an educated critical mass to deal with logic modelling and process modelling where logic challenges can be many and quite challenging. Also, there were no strong messages coming from the research community detailing the clear proven advantages if such technology were to be adopted.

3.2. WSMO/WSML/WSMX

WSMO (Web Services Modelling Ontology) [17] is a fully-fledged framework for Semantic Web Services. It appeared chronologically in 2005 after OWL/OWL-S, and was intended as a novel approach to address by design several limitations exhibited by OWL. The WSMO initiative provides a complete framework enhancing syntactic description of Web Services with semantic metadata. The WSMO project is an ongoing research and development initiative aiming to define a complete framework for Semantic Web Services and consists of three activities (i) WSMO,[9] which provides a formal specification of concepts for Semantic Web Services; (ii) WSML (Web Services Modelling Language),[10] which defines the language for representing WSMO concepts; (iii) WSMX[11] (Web Services Execution Environment), which defines and provides reference implementation allowing the execution of Semantic Web Services.

As illustrated in Fig. 4, WSMO was designed around concepts: Ontologies, Goals, Web Services and Mediators. As defined by [17], Goals provide a means to characterise user requests in terms of functional and non-functional requirements. For the former, a standard notion of pre- and post-conditions has been chosen and the latter provides a predefined ontology of generic properties. Web Service descriptions enrich this with an interface definition that defines service access patterns (its choreography) as well as the means to express services composed from other services (its orchestration).

The WSML language is based upon different logical formalisms: description logics, first-order logic and logic programming. WSML aims to provide a means to formally describe all the elements defined in WSMO. Some variants of WSML correspond to different levels of logical expressiveness/logic formalisms, and the use of different language paradigms. These variants are: WSML-Core, WSML-DL (description logic), WSML-Flight, WSML-Rule and WSML-Full. WSML-Core corresponds to the intersection of description logic and Horn logic. WSML is specified in terms of a normative human-readable syntax. The logic constraints that are declared in WSML goals or capabilities are interpreted and reasoned upon within WSMX by a dedicated component called

[8] It is important to note that RDF is not a subset of XML, rather that RDF/XML is only one of the multiple serialisation formats that can be used to represent RDF data.

[9] http://www.wsmo.org/.

[10] http://www.wsmo.org/wsml/.

[11] http://www.wsmx.org/.

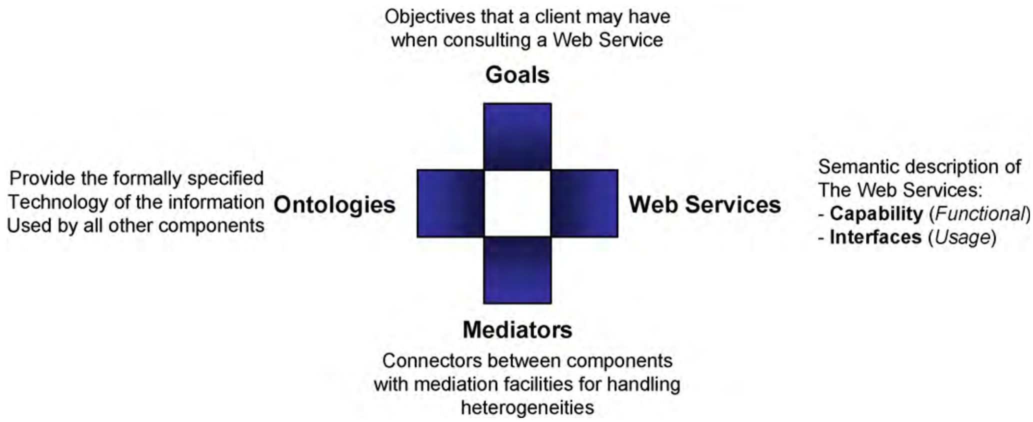

Fig. 4. WSMO top-level concepts.

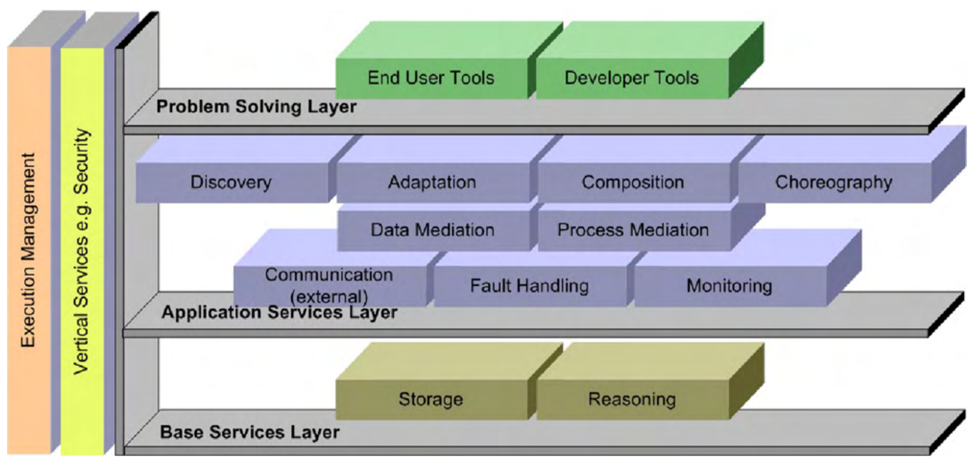

Fig. 5. WSMX reference architecture.

a 'reasoner'. Some of the reasoners available on the market include KAON2[12] [22], MINS[13] and IRIS [23].

As a conceptual and reference model, the WSMO framework provides the high-level concepts that are used in its WSMX reference implementation for Semantic Web Services. WSMX is an execution environment that supports dynamic discovery, selection, mediation, invocation and interoperation of Semantic Web Services, and provides a reference implementation for a service-oriented architecture that uses semantic annotation of all its major elements. The WSMX specification is currently being developed further through OASIS as the Semantic Execution Environment (SEE). In Fig. 5, we present the WSMX architecture and its most important components. WSMX is a framework for both Web Service providers and requesters. As a provider, one may register a service using WSMX in order to make it available to consumers and, as a requester, one can find the Web Services that suits one's needs and then invoke them in a transparent, secure and reliable way.

The platform shown in Fig. 5 offers the possibility to execute simple or complex operations with Web Services: discovery, invocation, composition (choreography and orchestration), mediation (where needed), reasoning over WS ontologies, and finally, results retrieval. The WSMX platform allows interoperability between WSML-described Web Services and other service definitions (e.g. via OWL) with the help of the data and process mediation components. Also, the platform allows various composition scenarios to be executed and for its components to be selectively used via an execution semantics process file that can be loaded into the WSMX core. With this platform, the semantic 'offering' of WSML and WSMO is completed and forms a full package for anyone who wants to use an open-source technology and have an ontology modelling methodology. This combination of a language (WSML) and an execution environment can be customised for various execution scenarios that can range across industries—tourism, banking, procurement and many others. Fig. 6.

To conclude, the SWS initiative has proved that it is possible to semi-automate processes and execute distributed operations over the Internet, while having a flexible and inferable way of interpreting data and behaviours attached to the services made available for machines. Various prototypes in e-tourism, e-banking, geospatial applications and business processes have been built in EU projects such as DIP [24], SWING [25] and SUPER [26], revealing that it is possible to build rather complex semi-automated applications and at the same time disclose the challenges of implementing such solutions on the market: the complexity of the technologies, the demand for highly-skilled knowledge workers, as well as the need for getting out on the market 'at the right time'. Overall, the combinations of the OWL/WSML ontologies, used for describing the interactions of Web Services with the help of inference provided by reasoners, is one of the most advanced applied research trials to date. The goals have been extremely ambitious and the prototypes delivered to date

12 http://kaon2.semanticweb.org/.
13 http://dev1.deri.at/mins.

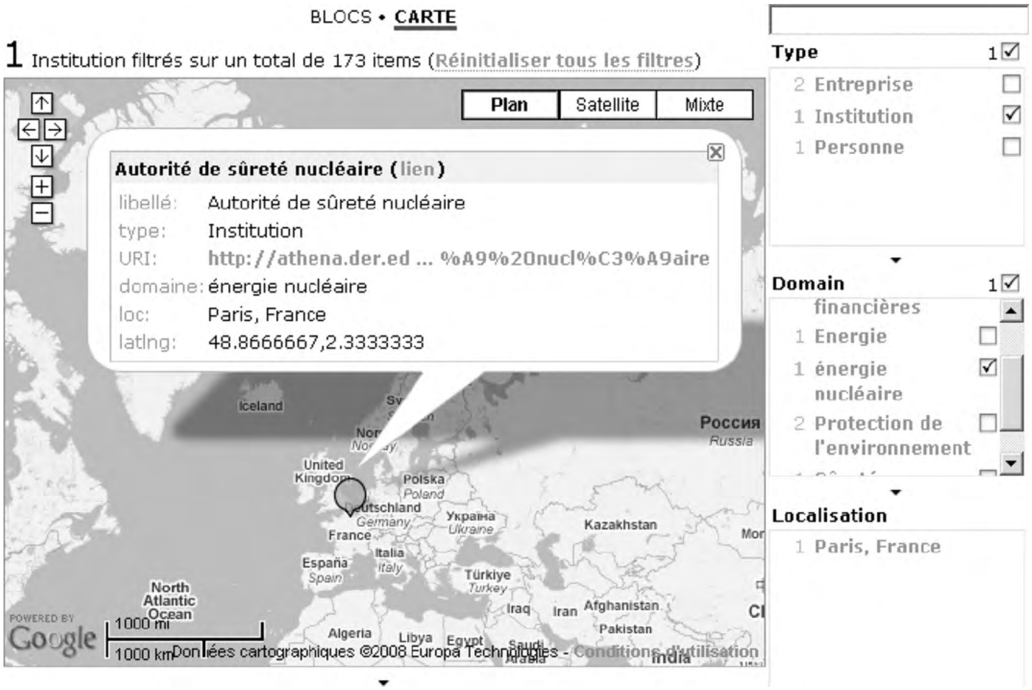

Fig. 6. Mashup of organisational data from a Semantic Wiki using Google Maps, GeoNames and Exhibit.

have proven the feasibility of such technical solutions, given that the right moment for them to be applied has been carefully chosen.

Sections 2 and 3 have presented the foundations of the Semantic Web. In the next three sections we look into each of the three applications domains outlined earlier—Knowledge Management, Enterprise Application Integration and E-Commerce. We begin with some Knowledge Management Applications. As Knowledge Management is a wide topic, especially in organisations, we will mainly focus on the use of collaborative software for Knowledge Management and in particular discuss their improvements thanks to Semantic Web technologies.

4. Knowledge Management and the Semantic Web

"Enterprise 2.0" [27] describes a next generation of Knowledge Management and collaboration tools being used in organisations, similar to how Web 2.0 is being used to describe a second-generation of web-based communities and hosted media services. It is defined as "the use of emergent social software platforms within companies, or between companies and their partners or customers".[14] For instance, blogs can be used to ease information sharing within organisations [28], wikis can be efficient platforms for collaborative document editing, from project management to software development [29], while microblogging can be used as a means to enhance real-time conversations and questions and answers between employees. Enterprise 2.0 is a new way to share and manage knowledge in industry, focusing on the "we are smarter than me"[15] meme [30] in order to enable collective intelligence in organisations.

In the same way that the Semantic Web refers to an extension of the Web providing structured and machine-readable content, Semantic Enterprise 2.0 refers to the extension of these tools using semantic technologies, e.g. semantic wikis or semantically-enhanced blogging within organisations [31]. Most of the tools used in these enterprise settings can be efficiently enhanced with Semantic Web technologies. For example, semantic blogging [32] systems such as semiBlog [33] allow one to embed metadata within blog posts where more structure is required, e.g. for meeting minutes or for project deliverable descriptions. Similarly, Haystack allows us to embed structured information and link to other resources when blogging [34]. More focused on real-time exchange in organisations, microblogging can also be enhanced with rich and structured metadata for similar purposes [35]. Semantic tagging capabilities, including the Tag Ontology [36], SCOT [37] and MOAT [38], can also help to integrate tagged data from different sources, the latter providing the ability to give meaning to users' tags using ontology instances, enabling enhanced information retrieval over tagged data [39].

One of the major areas for Semantic Web technologies and collaborative Knowledge Management in organisations is the use of semantic wikis. Semantic wikis enhance usual wiki functionalities by providing semantic annotations regarding the structure of the wiki, its content, or sometimes both. In his presentation on "The Relationship Between Web 2.0 and the Semantic Web",[16] Mark Greaves, previously at DARPA, said that semantic wikis are a promising answer to various issues associated with semantic authoring by professionals, by reducing the investment of time required for training on an annotation tool and by providing incentives required to providing semantic markup (attribution, visibility and reuse by others). Semantic wikis such as the popular Semantic MediaWiki extension [40] for the MediaWiki platform are also being repackaged for commercial use by companies like Centiare (now MyWikiBiz) or Ontoprise. Various semantic wiki applications exist, enabling the collaborative construction of structured knowledge and providing advanced querying or browsing interfaces, some of them also including reasoning capabilities. As well as Semantic MediaWiki, some interesting semantic wiki applications of note include: IkeWiki [41], OntoWiki [42], KiWi [43], and UfoWiki [44]. In [31,44], the use of a semantic wiki at EDF (Électricité de France, the main electricity company in France) is described, focusing on how ontologies are used to

[14] http://andrewmcafee.org/blog/?p=76.
[15] http://wearesmarter.org.
[16] http://tinyurl.com/markgreaves.

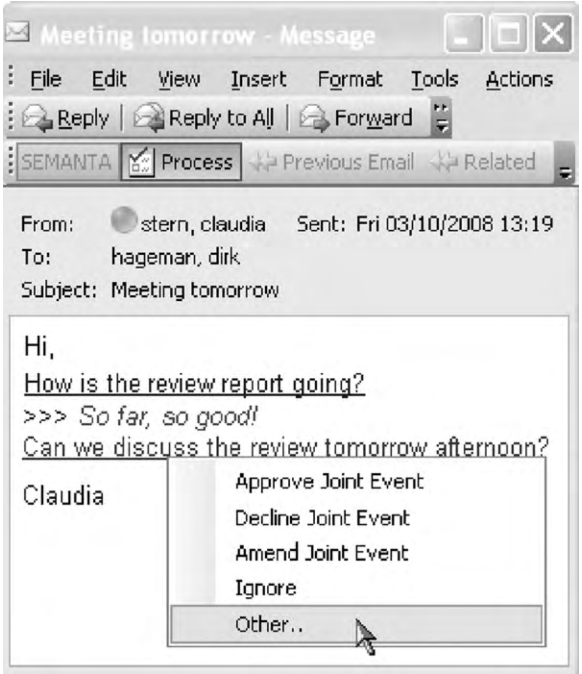

Fig. 7. Tracking actions in Semanta.

provide advanced means to identify relevant companies in collaboratively-built knowledge. It is also demonstrated how internal wikis can be enhanced from structured information from the Web, especially via the Linking Open Data project [45] to provide geolocation mashups combining internal wiki information and data from the Web (as depicted below). In [46], another use case regarding Semantic MediaWiki is described, focusing on the deployment of the application in a network of small- to medium-sized enterprises (SMEs).

4.1. Personal Knowledge Management

Semantic Web technologies in organisations can be used not only to enhance information published on Intranets, but also to improve personal information management and direct exchanges between employees. The Semantic Desktop initiative, via the Nepomuk project [47], currently deployed in KDE environments, uses dedicated ontologies to improve integration between desktop applications such as calendar, note taking, and so on. Projects within the Semantic Desktop initiative include Semanta [48], a semantic e-mail client that can identify various discourse patterns in e-mails such as actions, and then enhance the usual workflow of communication between knowledge workers, as depicted in Fig. 7, in which someone can directly accept or deny a meeting proposal that has been identified in the e-mail message. Other tools include SemNote [49], a semantic notetaking application, enabling the ability to link a note to any content on the desktop, providing integration between tools originally designed as independent data silos.

4.2. Public applications for corporate semantic Knowledge Management

Companies can also use public Web tools for Knowledge Management, especially when these tools allow private areas to be created. Many tools are now emerging that makes use of semantic representations for more efficient Knowledge Management and retrieval. Among these services, Radar Networks' flagship service is a semantic social software product called 'Twine' [50] that allows people to share what they know. It can be thought of as a knowledge networking application in that it allows users to share, organise, and find information with people they trust. People create and join 'twines' (community containers) around certain topics of interest, and items (documents, bookmarks, media files, etc.) are posted to these twines through a variety of methods. Content can be tagged by users when it is added, but Twine also automatically extracts relevant tags and entities from the content (people, places, organisations, etc.) Unlike many public communal sites, more than half of the data and activities in Twine are private (60%), and privacy/permission controls are deeply integrated into the Twine data structures. Another interesting service is Faviki [51] that allows people to annotate their content using tags based on DBpedia (the RDF version of Wikipedia) identifiers, enabling better ways to link tagged content together. In addition, Faviki exposes its data in RDF using CommonTag,[17] making integration easier with other services. Another San Francisco-based company, Metaweb, has created the open collaborative knowledge database Freebase, termed a "massive collaboratively-edited database of cross-linked data" powered by semantic technologies. Freebase organises its data and categories of data in ontology-like structures called "Freebase Types", based on a graph model. Any user can create and modify their own types and associated properties, and these can be promoted for adoption by administrators of the relevant domains that the type belongs to. Both Metaweb and Radar Networks have indicated that their services may later be repackaged for organisational use, thereby allowing proprietary or commercial data to be stored and accessed using semantic technologies [52]. As mentioned earlier, Knowledge Management is a relatively low risk but also low benefit area for IT investment in industry. We now turn our attention to an area where the benefits of Semantic Web can be higher, but with that comes a higher risk in terms of successful implementation—enterprise data integration within industry.

5. Enterprise data integration using semantic web technologies

By nature, Semantic Web technologies are a great candidate for enterprise data integration and middleware services [52]. By mapping data contained in heterogeneous applications to common representation layers using RDF(S)/OWL, they unify heterogenous data structures in a meaningful way. Among the various systems using Semantic Web technologies for data integration, Ontobroker [53] (now available as a commercial product from Ontoprise)[18] uses ontologies to extract and integrate information from distributed and heterogeneous semi-structured data sources, and provides a single query interface so that the user sees these distributed applications as a single system. SCORE – Semantic Content Organisation and Retrieval Engine [54] – uses Semantic Web technologies to integrate external sources inside management systems, using ontologies to manage heterogeneous data and using data extraction process from these various sources. In [55], the authors propose an ontology-based Knowledge Management System, aligning several internal data sources (databases, directories) via a mediation layer and aligning different local ontologies in the system.

Since lots of knowledge in organisations is stored inside relational databases management systems, another relevant area for enterprise data integration is also how this existing data can be integrated with Semantic Web applications. Recent efforts have provided various tools to translate existing databases to RDF, or to directly use SPARQL to query RDBMS. These include Triplify and D2RQ [56], and work regarding integrations between RDBMS and

[17] http://commontag.org.
[18] http://www.ontoprise.de/.

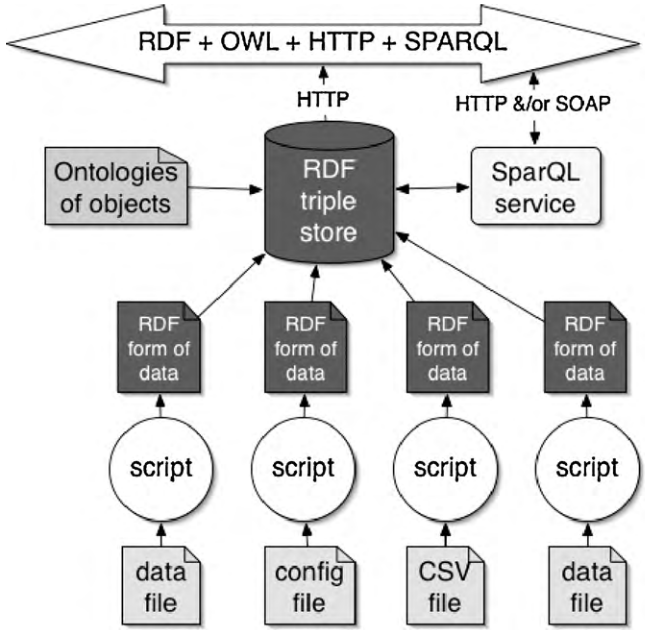

Fig. 8. The RDF Bus architecture.

RDF is to be tackled in the W3C by the RDB2RDF Working Group.[19] Recently, the W3C RDF2RDB Incubator Group[20] published a survey on current approaches regarding these mappings [57], as well as a report defining future research and standardisation steps on the topic [58]. Another means to bridge this gap between the RDF and XML worlds is XSPARQL [59], while GRDDL (Gleaning Resource Descriptions from Dialects of Languages) can be used to convert any XML document to RDF [60]. Moreover, it is worth mentioning that various use cases have successfully reported the use of data integration using Semantic Web technologies systems, within companies ranging from NASA [39], Lilly [61], British Telecom [62], and a recent W3C workshop on Semantic Web technologies in Energy Industries[21] also showed interest from companies such as Chevron in the area.[22]

5.1. Data integration in collaborative environments

We previously mentioned how Enterprise 2.0 services can be independently enhanced using Semantic Web technologies, but there is often a need to integrate various sources (blogs, wikis and RSS feeds) together to have a global object-centric [63] view of a given topic. In order to define a mediation layer to interlink these social applications, one can rely on lightweight vocabularies such as SIOC [64] and FOAF [9] to unify the content from these applications. These applications can interact using an architecture similar to the RDF Bus, introduced by [65] and described in [66], and as depicted in the services shown in Fig. 8.

Collaborative work environments (CWEs) used for collaboration between organisations often act as data silos since they are isolated systems that cannot communicate with the rest of the world. SIOC has been adapted in the ECOSPACE integrated project for use in CWEs as an interoperable format to share workspaces and documents between different CWE systems, e.g. BSCW and Business Collaborator [67]. Concepts that exist in the CWE domain were mapped to the SIOC ontology (e.g. document, folder, user and post). CWEs can then access data from heterogeneous remote CWEs within the application itself.[23] A user in BSCW or BC can create a new "shadow" folder which connects to a remote folder in a remote (heterogeneous) CWE, thereby providing transparency and the same look-and-feel as native folders[24] with much of the same essential functionality.

We now turn our attention to the final and arguably the most risk-intensive application domain for Semantic Web technologies—e-commerce. This domain is also arguably the domain with the highest potential rewards for industry from successful implementation. This final domain, e-commerce, builds on both the Semantic Web and Semantic Web Services discussed in Sections 2 and 3.

6. E-Commerce

The term e-commerce as used in this article relates to the applications that facilitate business-to-consumer and business-to-business transactions using elements of Semantic Web technologies and Semantic Web Services discussed earlier, to improve and enhance value-adding processes. Current transactions depend to a large extent on human interaction between co-operating processes and/or on vendor-specific software that utilises various standards available in this domain. Semantic Web technologies and services can provide mechanisms that integrate applications and services in an extended and virtual business environment. Virtual enterprises can be envisaged that will support configuration and reconfiguration through semi-automatic service discovery and mediation. Major information transactions from functions such as design, production, planning, distribution, supply chain, recycling and the innovation process can be codified into an ontology and made machine processable (see Fig. 9).

The use of a loosely-integrated virtual enterprise-based framework holds the potential of adapting to changing market demands. Its success requires reliable and large-scale interoperation among trading partners via a Semantic Web of trading partners' services whose properties, capabilities, and interfaces are encoded in an unambiguous as well as computer-understandable form. The ability to search for and quickly find the small piece of information needed from the huge amount of information available has crucial importance. To overcome current bottlenecks in business-to-business (B2B) electronic commerce, we need intelligent solutions for mechanising the process of structuring, standardising, aligning and personalising data. Interoperability represents a major challenge for current IT solutions in dealing with the information layer of business-to-business (B2B) integration that may only be overcome by means of an ontology-based mediation. This allows organisations with different data standards to exchange information seamlessly without having to change their proprietary data schemas. The Semantic Web can make e-commerce interactions more flexible and automated by standardising ontologies, message content, and message protocols. Exploration of Semantic Web technologies in particular ontologies, grid technologies and Semantic Web Services promotes the development of modern manufacturing technology. Grid applications need resources and services to be discovered quickly and efficiently. The greatest obstacle however, lies in the difficulties in appropriately describing resources and services. To easily find the needed services, an ontology is used to describe resources and services. Recent work in Artificial Intelligence (AI) is exploring the use of formal ontologies as a way of specifying content-specific agreements for the sharing and reuse of knowledge among software entities. Formal ontologies are viewed as designed

[19] http://www.w3.org/2009/03/rdb2rdf-charter.
[20] http://www.w3.org/2005/Incubator/rdb2rdf/.
[21] http://www.w3.org/2008/12/ogws-cfp.
[22] http://www.w3.org/2008/07/ogws-report.html.
[23] http://tinyurl.com/siocecospace2.
[24] http://tinyurl.com/siocecospace2.

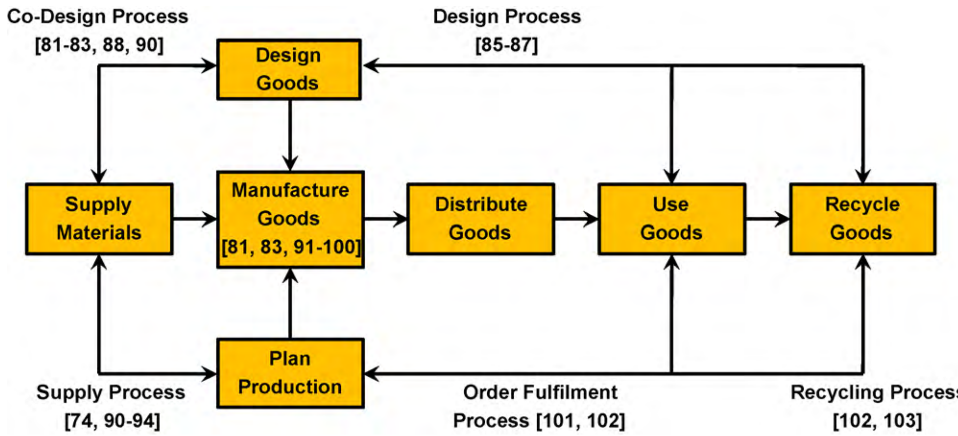

Fig. 9. Functions and processes in industry.

artefacts, formulated for specific purposes and evaluated against objective design criteria.

The vision of the Semantic Web is to have many routine transactions between all of these functions and processes across the extended enterprise become machine processable, thus minimising human interaction and the need for expensive bespoke programming. The scale of processes and applications across e-commerce is significant. The spectrum of applications is extremely large and to name a few: corporate portals and knowledge management, e-commerce, e-work, e-business, healthcare, e-government, natural language understanding and automated translation, information search, data and services integration, social networks and collaborative filtering, knowledge mining, business intelligence and so on.

Fig. 9 illustrates a number of generic functions in a typical industrial organisation that deploy computer technology to support its business processes, in particular B2B and B2C communications. The figure also illustrates key processes that integrate specific functions, for example, the co-design process that integrates the design, supply and production functions. Other processes include the concurrent design process, the e-business process, the supply chain process, the order fulfillment process and the recycling process. There are a significant number of research initiatives for creating ontologies and associated services and applications across these and other functions and processes. Fig. 9 also illustrates some published research on the key research thrusts that are emerging. Not included in this figure is the large number of surveys completed into industrial-based Semantic Web applications [68–73]. Also, not included in this figure are a number of research papers that are emerging on the application of Semantic Web technologies in specific application domains that vary from organic food [74], telecommunications [75,76], and petroleum [77,78] to the pharmaceutical industry [79] and healthcare [80]. We will now briefly examine some of the research underway in the core business processes illustrated.

6.1. Co-design and design processes

The co-design process takes place between the design function and two major stakeholders—suppliers and customers. Qiu [81] looks at service-based collaborative design architectures to allow multi-disciplinary experts to cooperate closely in the whole product development lifecycle. Lee [82] looks at distributed product development architecture for engineering collaborations across ubiquitous virtual enterprises. Finally, Lin [83] looks at a manufacturing system engineering ontology for semantic interoperability across extended project teams. The principle objectives of this research are to find an efficient method through which partnerships may be managed on demand and to facilitate collaborative design. The key solution emerging are the use of Web Services for the discovery of common processes, a common ontology that promotes integration between dissimilar information systems and applications codified in machine-readable languages such as RDF and OWL. Various service models have emerged which are able to represent manufacturing behaviour, facilitate the complex communication required for collaborative process management, and help companies collaborate on an optimal solution for design and co-design of products and processes. Other ontology-based knowledge schemes facilitate communication and information exchange in inter-enterprise, multi-disciplinary engineering design teams and are encoded in standard Semantic Web languages. The approach which is common across a wide range of research in this area focuses on how to support information autonomy that allows the individual team members to keep their own preferred languages or information models rather than requiring them all to adopt a standardised terminology. The use of semantics extends beyond product data and includes the modelling of skills and expertise among design team members. The complexity of design data increases substantially in extended enterprise environments comprising of a large number of co-designing organisations. Recent advances in grid, Semantic Web, P2P and Web Services have revolutionised the way we communicate and collaborate and also make it possible to integrate desktop data. Undertaking the intersection fulfillment of these technologies on an enterprise desktop enables new integrative and collaborative use of not only organisational data but also distributed, autonomous personal data on the Web by fully leveraging potentially-useful information sources on many desktops. Collaborative product commerce is carried out concurrently with the identification of potential manufacturing partners based on the design requirements and manufacturing constraints. The innovation management process includes research thrusts in the fields of creativity and project management. Specific research in this area is currently limited but some initial thrusts are emerging regarding how Semantic Web technologies can be leveraged to improve the innovation management process itself [84] and how it can be used in project management [85–87]. Lin's work on extended project teams is also relevant here [88]. A recent effort in the domain of e-commerce and Semantic Web technologies is the GoodRelations vocabulary.[25] The aim of this vocabulary is to let vendors and business annotate their content with machine-readable descriptions of the

[25] http://www.heppnetz.de/projects/goodrelations/.

products they are selling (price, quantity, etc.) as well as information about the shop itself (opening time, allowed payment types, etc.). In order to facilitate the annotation process, different tools have been provided, such as the GoodRelations annotator.[26] Thanks to these annotations, if widely deployed on the Web (BestBuy has already released details of more than 1000 stores in the US using this model), one can run queries to find, for example, all places in his or her locality that are open on weekends and that sell CD players for less than $20.

6.2. Supply process

The supply chain process encompasses relationships between organisations and their suppliers, and typically involves business-to-business digital transactions. There is a significant amount of research in this area, most of which has been discussed earlier in the context of Semantic Web Services. Important research topics that include Semantic Web technologies include collaborative commerce (c-commerce) [89], product data integration [90], partnership outsourcing [91] and matchmaking [92]. Other important topics include the concept of dynamically building virtual enterprises that come together to co-design and co-manufacture goods and services. Moitra explores the creation of adaptive enterprises [93] whereas Shen looks at ways to configure networks of manufacturing competences and capabilities [94]. Web Services have emerged as a promising enabling technology for BPM in support of c-commerce. Web Services offer effective and standard-based means to improve interoperability among different software applications over Internet protocols. Web Services can enhance business process coordination and provide rapid reconfigurability in order to evolve and adapt to emerging issues such as mass customisation. This challenge is increased as new types of processes and components are introduced, as existing components are expected to interact with the novel entities but have no previous knowledge on how to collaborate. The capabilities of Semantic Web Services for performing automatic service discovery, selection, composition and invocation enable manufacturing systems to self-orchestrate without the need for manual configuration, and without the need for concentrating logic in centralised systems.

6.3. Production process

The production process focuses attention onto the shop floor of manufacturing and in particular into how Semantic Web technologies and sensor technologies can be put to use to improve costs. There are a number of research thrusts emerging in this domain including resource discovery [95], factory automation [91], and service-oriented manufacturing resource planning [96]. Various ontologies are emerging to help manage the various common concepts used in production [94,97,98]. Some examples include an array of architecture patterns for creating distributed message brokers frameworks, focusing mainly on globally-distributed federations and locally-distributed clusters. A unified architecture is subsequently presented that leverages the different patterns by combining federated frameworks with locally-distributed clusters into a unified set of architecture elements and interactions. The service-oriented approach is also used to dynamically discover resources and automatically invoke the (re)configuration and messaging services. The services are enriched with semantics in order to facilitate automatic discovery and selection of services using the Semantic Web Services formal ontology. One of the significant challenges is that of providing rapid reconfigurability in order to evolve and adapt to mass customisation. This challenge is aggravated if new types of processes and components are introduced, as existing components are expected to interact with the novel entities but have no previous knowledge on how to collaborate. This statement not only applies to innovative processes and devices, but is also due to the impossibility to incorporate knowledge in a single device about all types of available system components.

6.4. Order fulfillment process

The order fulfillment process looks at the way customers are linked to the manufacturing process through the order cycle and customer relationship management processes. This process and indeed others mentioned above are currently dominated by large manufacturing resource planning and customer relationship management systems produced by the large software houses including SAP, Baan, Microsoft and Oracle. Semantic Web technologies are slow to penetrate these traditionally strong software sectors that typically integrate their own systems before concentrating on integration with third-party software. Outside of the mainstream software vendors, a number of research thrusts are emerging including ontologies for e-business [99], and fuzzy-ontology generation for help desk support [100]. Service Orientated Architectures (SOA) based on Semantic Web Services will allow organisations to enhance interoperability and encourage reuse of components and interfaces. The application of semantic descriptions to services is frequently advocated in research with the aim of further improving SOA and enabling scalability. Components forming part of an SOA can be descried semantically in term of commonly-understood data and process ontologies. Various technical difficulties and ontological issues arise, such as distributed computing, application integration, and distributed product information sharing that can impede the collaborative product development. One issue is concerned with product information sharing and synchronisation across virtual enterprises. Another is concerned with federations of product development services over the collaborative process. The third is related to the engineering context management using the Semantic Web for providing more human-oriented collaborations services. The service orientation is achieved via virtualisation in which everything, including machines, equipments, devices, various data sources, applications, and processes, are virtualised as standard Web Services.

6.5. Recycling process

The product life cycle now extends to end of life. This has an impact on the way organisations design their products, i.e. design for the environment and also the way products can be efficiently recycled once they reach their end of life. Surprising little research is ongoing in this area with respect to Semantic Web technologies. Two interesting research thrusts include a semantically-enabled compliance tool for environment-focused design [101] and tools for assisting the design process to discover more environmentally-friendly materials used in products [102].

6.6. Trends

We end this review of e-commerce by citing a major report on the application of semantic technologies in industry. The report entitled 'Semantic Wave Report—Industry Roadmap to Web 3.0' was released in 2008 and reports on the potential market impact of the Semantic Web and how it will add value to existing information systems [103]. It charts the development of the Internet from the current Web 2.0 towards Web 3.0 and visualises semantic technologies making the Internet and Web even more

[26] http://www.ebusiness-unibw.org/tools/goodrelations-annotator/.

Table 1
Value proposition of the Semantic Web (source: Semantic Wave 2008 Report [103]).

Challenges	Semantic capabilities	Value drivers
1. Development: Complexity, labour-intensity, solution time, cost, risk	Semantic modelling is business rather than IT centric, flexible, less resource intense, and handles complex development faster	Semantic automation of "*business need-to-capability-to-simulate-to-test-to-deploy-to-execute*" development paradigm
2. Infrastructure: Net-centricity, scalability; resource, device, system, information source, communication intensity	Semantic enablement and orchestration of transport, storage, and computing resources, IPv6, SOA, WS, BPM, EAI, EII, grid, P2P, security, mobility, system-of-systems	In the semantic wave, infrastructure scale, complexity, and security become unmanageable without semantic solutions
3. Information: Semantic interoperability of information formats, sources, processes, and standards; search relevance, use context	Composite applications (information and applications in context powered by semantic models), semantic search, semantic collaboration, semantic portals	Semantic interoperability, semantic search, semantic social computing, and composite applications and collaborative knowledge management become "killer apps"
4. Knowledge: Knowledge automation, complex reasoning, knowledge commerce	Executable domain knowledge-enabled authoring, research, simulation, science, design, logistics, engineering, virtual manufacturing, policy and decision support	Executable knowledge assets enable new concepts of operation, super-productive knowledge work, enterprise knowledge superiority, and new intellectual property
5. Behaviour: Systems that know what they are doing	Robust adaptive, autonomic, autonomous system behaviours, cognitive agents, robots, games, devices, and systems that know, learn, and reason as humans do	Semantic wave systems learn and reason as humans do, using large knowledge bases and reasoning with uncertainty and values as well as logic

ubiquitous as it begins to impact all of the sciences and utilises a wide variety of ambient intelligent devices and sensors. The report created a table of how semantic technologies will add value, see Table 1. Value is realised across five interrelated phases—development, infrastructures, information, knowledge and behaviour.

The report concludes by describing a number of trends that will bring the Semantic Web to the next level of diffusion. These trends include: intelligent user interfaces; semantic social computing; semantic applications across a wide variety of industries; semantic infrastructure to support Semantic Web technologies and; semantic development environments. It cites over 270 commercial organisations that are currently offering value to customers with products built on Semantic Web technologies.

7. Conclusions

In this article, we have reviewed the use of Semantic Web technologies and Semantic Web Services in various applications in the industrial domain and in particular Knowledge Management, Enterprise Application Integration and E-Commerce. The Semantic Web is now an active topic of research among leading research groups around the world. It is also the core enabling technology in an increasingly wide range of commercial applications. Many of these are breakthrough applications offering enhanced functionality beyond conventional applications. As some of these Semantic Web applications begin to make a breakthrough and attract attention from more routine and traditional solution providers, it can be anticipated that there will be a wider adoption of Semantic Web technologies among the programming community, thereby leading to a potentially exponential increase in the adoption of the Semantic Web and associated new applications and IT solutions. The Semantic Web has matured as an enabling technology but is still in its infancy regarding industrial applications. Through this review, we have demonstrated how the Semantic Web has moved from being an academic endeavour (where the various languages, logics and representation mechanisms required for it to succeed are now in place and have been standardised within W3C during the past few years) to a commonly-deployed industrial technology (with a wide range of applications for small, medium and large enterprises as well as dedicated solution vendors). Various challenges are still to be taken into account, such as the dynamics of some social information systems (e.g. semantic microblogging) as well as integration with sensor data, that can be used for example to track purchases and deliveries more precisely. It is clear from the number of new journals and conferences dedicated to the Semantic Web that it has become a main stream branch of computer science. What remains to be demonstrated is whether it has the potential to become the next breakthrough technology that will lead industry to the next level of value-creating solutions for customers.

Acknowledgment

This work was supported by Science Foundation Ireland under grant number SFI/08/CE/I1380 (Líon 2).

References

[1] J. Markoff, Entrepreneurs see a web guided by common sense, New York Times (2006).
[2] J. Cardoso, The Semantic Web vision: where are we? IEEE Intelligent Systems 22 (5) (2007) 84–88.
[3] P. Mika, Microsearch: an interface for semantic search, in: Proceedings of the SemSearch, 2008, pp. 79–88.
[4] T. Berners-Lee, J.A. Hendler, O. Lassila, The Semantic Web, Scientific American 284 (May (5)) (2001) 34–43.
[5] T.R. Gruber, A translation approach to portable ontology specifications, Knowledge Acquisition 5 (1993) 199–220.
[6] D. Bearman, E. Miller, G. Rust, J. Trant, S. Weibel, A common model to support interoperable metadata, D-Lib Magazine (January) (1999) http://www.dlib.org/dlib/january99/bearman/01bearman.html.
[7] O. Lassila, R. Swick, Resource description framework (RDF) model and syntax specification, W3C Recommendation (February) (1999).
[8] D. Brickley, R. Guha, RDF vocabulary description language 1.0: RDF schema, W3C Recommendation (February) (2004) http://www.w3.org/TR/rdf-schema/.
[9] D. Brickley, L. Miller, FOAF Vocabulary Specification, 2004 http://xmlns.com/foaf/spec.
[10] C. Bizer, R. Cyganiak, T. Heath, How to Publish Linked Data on the Web, Tutorial, 2007 http://www4.wiwiss.fu-berlin.de/bizer/pub/LinkedDataTutorial.
[11] D. McGuinness, F. van Harmelen, OWL Web Ontology Language Overview, 2004 http://www.w3.org/TR/owl-features/.
[12] E. Prud'hommeaux, A. Seaborne, "SPARQL Query Language for RDF" and "SPARQL Protocol for RDF", W3C Recommendations, January, 2008, http://www.w3.org/TR/rdf-sparql-query/ and http://www.w3.org/TR/rdf-sparql-protocol/.
[13] K. Kjernsmo, A. Passant, SPARQL New Features and Rationale, 2009 http://www.w3.org/TR/sparql-features/.
[14] E. Miller, Reuse, repurpose, remix, in: Proceedings of the Semantic Technologies Conference 2008, San Jose, CA, May, 2008.
[15] T. Berners-Lee, Linked data, W3C Design Issues. http://www.w3.org/DesignIssues/LinkedData.html.
[16] D. Martin, M. Burstein, J. Hobbs, O. Lassila, D. McDermott, S. McIlraith, S. Narayanan, M. Paolucci, B. Parsia, T. Payne, E. Sirin, N. Srinivasan, K. Sycara, OWL-S—Semantic Markup for Web Services, 2004 http://www.w3.org/Submission/OWL-S/.
[17] D. Roman, H. Lausen, U. Keller, Web Service Modelling Ontology, 2005 http://www.wsmo.org/TR/d2/v1.3/.

[18] F. Baader, D. Calvanese, D.L. McGuinness, D. Nardi, P.F. Patel-Schneider, The Description Logic Handbook: Theory, Implementation, Applications, Cambridge University Press, Cambridge, UK, 2003.
[19] F. Wiedijk, J. Zwanenburg, First Order Logic with Domain Conditions, Theorem Proving in Higher Order Logics, Springer, Berlin, 2003.
[20] J.W. Lloyd, Foundations of Logic Programming, 2nd ed., Springer, 1987.
[21] OWL-S W3C Member Submission. http://www.w3.org/Submission/OWL-S/.
[22] B. Motic, Reasoning in description logics using resolution and deductive databases, Ph.D. Thesis, University of Karlsruhe, Karlsruhe, Germany, January, 2006.
[23] IRIS reasoner. http://www.iris-reasoner.org/.
[24] DIP project. http://dip.semanticweb.org/index.html.
[25] SWING project. http://www.swing-project.org/.
[26] SUPER project. http://www.ip-super.org/.
[27] A. McAfee, Enterprise 2.0: the dawn of emergent collaboration, MIT Sloan Management Review 47 (3) (2006) 21–28.
[28] P. Kolari, T. Finin, K. Lyons, Y. Yesha, Y. Yesha, S. Perelgut, J. Hawkins, On the structure, properties and utility of internal corporate blogs, in: Proceedings of the 1st International Conference on Weblogs and Social Media (ICWSM 2007), Boulder, CO, 2007.
[29] A. Majchrzak, C. Wagner, D. Yates, Corporate Wiki users: results of a survey, in: Proceedings of the International Symposium on Wikis, 2006.
[30] B. Libert, J. Spector, D. Tapscott, We Are Smarter Than Me: How to Unleash the Power of Crowds in Your Business, Wharton School Publishing, 2007.
[31] A. Passant, Case study: enhancement and integration of corporate social software using the semantic web, in: Semantic Web Use Cases and Case Studies, W3C Semantic Web Education and Outreach Interest Group, 2008 http://www.w3.org/2001/sw/sweo/public/UseCases/EDF/.
[32] D.R. Karger, D. Quan, What would it mean to blog on the semantic web? in: Proceedings of the 3rd International Semantic Web Conference (ISWC 2004), Hiroshima, Japan, November, 2004.
[33] K. Möller, U. Bojārs, J.G. Breslin, Using semantics to enhance the blogging experience, in: Proceedings of the 3rd European Semantic Web Conference (ESWC' 06), LNCS vol. 4011, Budva, Montenegro, June, (2006), pp. 679–696.
[34] D. Quan, D. Huynh, D.R. Karger, Haystack: a platform for authoring end user semantic web applications, in: Proceedings of the Second International Semantic Web Conference (ISWC 2003), LNCS vol. 2870, Sanibel Island, FL, USA, (2003), pp. 738–753, ISBN: 3-540-20362-1.
[35] A. Passant, T. Hastrup, U. Bojars, J.G. Breslin, Microblogging: a semantic web and distributed approach, in: Proceedings of the 4th Workshop on Scripting for the Semantic Web at the 5th European Semantic Web Conference, Tenerife, Spain, June, 2008.
[36] R. Newmann, "The Tag Ontology". http://www.holygoat.co.uk/projects/tags/.
[37] H.L. Kim, S.K. Yang, J.G. Breslin, H.G. Kim, Simple algorithms for representing tag frequencies in the SCOT exporter, in: Proceedings of the IEEE/WIC/ACM International Conference on Intelligent Agent Technology, Fremont, CA, USA, (2007), pp. 536–539.
[38] A. Passant, P. Laublet, Meaning of a tag: a collaborative approach to bridge the gap between tagging and linked data, in: Proceedings of the 1st Workshop on Linked Data On the Web (LDOW 2008), Beijing, China, 2008.
[39] M. Grove, A. Schain, NASA's expertise location service powered by semantic web technologies, in: W3C Semantic Web Use Cases and Case Studies, 2008 http://www.w3.org/2001/sw/sweo/public/UseCases/Nasa/.
[40] M. Krötzsch, D. Vrandecic, M. Völkel, Semantic MediaWiki, in: Proceedings of the 5th International Semantic Web Conference (ISWC 2006), LNCS, vol. 4273, Georgia, USA, (2006), pp. 935–942.
[41] S. Schaffert, IkeWiki: a semantic Wiki for collaborative knowledge management, in: Proceedings of the 1st International Workshop on Semantic Technologies in Collaborative Applications (STICA '06), Manchester, UK, June, 2006.
[42] S. Auer, S. Dietzold, T. Riechert, OntoWiki—a tool for social, semantic collaboration, in: Proceedings of the 5th International Semantic Web Conference (ISWC 2006), LNCS vol. 4273, Athens, GA, USA, November, (2006), pp. 736–749.
[43] S. Schaffert, J. Eder, S. Grünwald, T. Kurz, M. Radulescu, KiWi—a platform for semantic social software, demonstration session, in: Proceedings of the 6th European Semantic Web Conference, Heraklion, Greece, June, (2009), pp. 888–892.
[44] A. Passant, P. Laublet, Combining structure and semantics for ontology-based corporate Wikis, in: Proceedings of the 11th International Conference on Business Information Systems (BIS 2008), Innsbruck, Austria, May, (2008), pp. 58–69.
[45] C. Bizer, T. Heath, T. Berners-Lee, Linked data—the story so far, in: T. Heath, M. Hepp, C. Bizer (Eds.), Special Issue on Linked Data, International Journal on Semantic Web and Information Systems, 2009.
[46] C. Christl, C. Ghidini, J. Guss, V. Pammer, S. Lindstaedt, P. Scheir, L. Serafini, M. Rospocher, Deploying semantic web technologies for work integrated learning in industry. A comparison: SME vs. large sized company, in: In Use Track, Proceedings of the 7th International Semantic Web Conference (ISWC 2008), 2008.
[47] A. Bernardi, S. Decker, L. van Elst, G. Grimnes, T. Groza, S. Handschuh, M. Jazayeri, C. Mesnage, K. Möller, G. Reif, M. Sintek, The social semantic desktop: a new paradigm towards deploying the semantic web on the desktop, in: J. Cardoso, M.D. Lytras (Eds.), Semantic Web Engineering in the Knowledge Society, IGI Global, 2008.
[48] S. Scerri, B. Davis, S. Handschuh, M. Hauswirth, Semanta—semantic email made easy, in: Proceedings of the 6th European Semantic Web Conference (ESWC 2009), Heraklion, Greece, 2009.
[49] L. Dragan, S. Handschuh, SemNotes—note-taking on the semantic desktop, in: Proceedings of the 6th European Semantic Web Conference (ESWC 2009), Poster Session, Heraklion, Greece, 2009.
[50] J. Wissner, N. Spivack, Case study: twine, in: W3C Semantic Web Use Cases and Case Studies, 2009 http://www.w3.org/2001/sw/sweo/public/UseCases/Twine/.
[51] V. Miličić, Case study: semantic tags, in: W3C Semantic Web Use Cases and Case Studies, 2008 http://www.w3.org/2001/sw/sweo/public/UseCases/Faviki/.
[52] G. Wiederhold, Mediators in the architecture of future information systems, IEEE Computer 25 (3) (1992) 38–49.
[53] S. Decker, M. Erdmann, D. Fensel, R. Studer, Ontobroker: ontology-based access to distributed and semi-structured information, in: Database Semantics: Semantic Issues in Multimedia Systems, 1999, pp. 351–369.
[54] A.P. Sheth, C. Bertram, D. Avant, B. Hammond, K. Kochut, Y.S. Warke, Managing semantic content for the web, IEEE Internet Computing 6 (4) (2002) 80–87.
[55] A. Maedche, B. Motik, L. Stojanovic, R. Studer, R. Volz, Ontologies for enterprise knowledge management, IEEE Intelligent Systems 18 (2) (2003) 26–33.
[56] C. Bizer, A. Seaborne, D2RQ—treating non-RDF databases as virtual RDF graphs, in: Proceedings of the 3rd International Semantic Web Conference (ISWC 2004), Poster Session, November, 2004.
[57] S. Sahoo, W. Halb, S. Hellmann, K. Idehen, T. Thibodeau Jr., S. Auer, J. Sequeda, A. Ezzat, A Survey of Current Approaches for Mapping of Relational Databases to RDF, W3C RDF2RDB XG, 2009 http://www.w3.org/2005/Incubator/rdb2rdf/RDB2RDF_SurveyReport.pdf.
[58] A. Malhotra, W3C RDB2RDF XG Report, 2009 http://www.w3.org/2005/Incubator/rdb2rdf/XGR/.
[59] W. Akhtar, J. Kopecky, T. Krennwallner, A. Polleres, XSPARQL: travelling between the XML and RDF worlds—and avoiding the XSLT pilgrimage, in: Proceedings of the 5th European Semantic Web Conference (ESWC 2008), 2008.
[60] D. Connolly, Gleaning resource descriptions from dialects of languages (GRDDL), W3C Recommendation (September) (2007) http://www.w3.org/TR/grddl/.
[61] S. Stephens, Prioritization of biological targets for drug discovery, in: W3C Semantic Web Use Cases and Case Studies, 2007 http://www.w3.org/2001/sw/sweo/public/UseCases/Lilly/.
[62] A. Duke, B2B integration with semantic mediation, in: W3C Semantic Web Use Cases and Case Studies, 2007 http://www.w3.org/2001/sw/sweo/public/UseCases/BT/.
[63] J.G. Breslin, S. Decker, The future of social networks on the internet: the need for semantics, IEEE Internet Computing 11 (November–December) (2007) 86–90.
[64] J.G. Breslin, A. Harth, U. Bojārs, S. Decker, Towards semantically-interlinked online communities, in: Proceedings of the 2nd European Semantic Web Conference (ESWC 2005), LNCS vol. 3532, Heraklion, Greece, May, (2005), pp. 500–514.
[65] T. Berners-Lee, Putting the web back in semantic web, in: Proceedings of the 4th International Semantic Web Conference, Wrap-Up Session and RuleML Kickoff Sessions, Galway, Ireland, 2005 http://www.w3.org/2005/Talks/1110-iswc-tbl/.
[66] A. Passant, Technologies du Web Semantique pour l'Entreprise 2.0 (Semantic Web Technologies for Enterprise 2.0), Ph.D. Thesis, 2009.
[67] U. Bojārs, J.G. Breslin, V. Peristeras, G. Tummarello, S. Decker, Interlinking the social web with semantics, IEEE Intelligent Systems 23 (May (3)) (2008).
[68] W. Abramowicz, K. Haniewicz, M. Kaczmarek, D. Zyskowski, E-Marketplace for semantic web services, in: Proceedings of the 6th International Conference on Service-Oriented Computing, Sydney, Australia, 2008.
[69] S. Boddy, Y. Rezgui, G.S. Cooper, M. Wetherill, Computer integrated construction: a review and proposals for future direction, Advances in Engineering Software 38 (10) (2007) 677–687.
[70] H. Chen, J. Ma, Y. Wang, Z. Wu, A survey on semantic e-science applications, Computing and Informatics 27 (1) (2008) 5–20.
[71] A. Léger, J. Heinecke, L.J.B. Nixon, P. Shvaiko, J. Charlet, P. Hobson, F. Goasdoué, The semantic web from an industry perspective, in: Proceedings of the 2nd Reasoning Web Summer School, 2006, pp. 232–268.
[72] C. Legner, B. Lebreton, Preface to the focus theme section: business interoperability; business interoperability research: present achievements and upcoming challenges, Electronic Markets 17 (3) (2007) 176–186.
[73] H. Schevers, R. Drogemuller, Converting the industry foundation classes to the web ontology language, in: Proceedings of the 1st International Conference on Semantics, Knowledge and Grid (SKG' 05), 2005.
[74] C. Brewster, AKTive food: semantic web based knowledge conduits for the organic food industry, in: Workshop on Semantic Web Case Studies and Best Practices for eBusiness (SWCASE' 05), Galway, Ireland, 2005.
[75] A. Duke, J. Davies, M. Richardson, N. Kings, A semantic service orientated architecture for the telecommunications industry, Intelligence in Communication Systems (2004) 236–245.
[76] F. Gandon, O. Corby, A. Giboin, N. Gronnier, C. Guigard, Graph-based inferences in a semantic web server for the cartography of competencies in a telecom valley, in: Proceedings of the 4th International Semantic Web Conference (ISWC 2005), Galway, Ireland, 2005.
[77] J.A. Gulla, S.L. Tomassen, D. Strasunskas, Semantic interoperability in the Norwegian petroleum industry, in: Proceedings of the 5th International Conference on Information Systems Technology and its Applications (ISTA 2006), 2006, pp. 81–94.
[78] D. Norheim, R. Fjellheim, AKSIO—active knowledge management in the petroleum industry, industry forum, in: Proceedings of the 3rd European Semantic Web Conference (ESWC' 06), Budva, Montenegro, 2006.
[79] V. Venkatasubramanian, C. Zhao, G. Joglekar, A. Jain, L. Hailemariam, P. Suresh, P. Akkisetty, K. Morris, G.V. Reklaitis, Ontological informatics infrastructure for pharmaceutical product development and manufacturing, Computers and Chemical Engineering 30 (10–12) (2006) 1482–1496.
[80] C. Wroe, Is semantic web technology ready for healthcare? Industry Forum, in: Proceedings of the 3rd European Semantic Web Conference (ESWC' 06), Budva, Montenegro, 2006.

[81] H.B. Qiu, X.Y. Shao, P.G. Li, L. Gao, An agent- and service-based collaborative design architecture under a dynamic integration environment, The International Journal of Advanced Manufacturing Technology 35 (1) (2007) 15–25.
[82] J. Lee, K. Kim, A distributed product development architecture for engineering collaborations across ubiquitous virtual enterprises, International Journal of Advanced Manufacturing Technology 33 (1) (2007) 59–70.
[83] H.K. Lin, J.A. Harding, A manufacturing system engineering ontology model on the semantic web for inter-enterprise collaboration, Computers in Industry 58 (5) (2007) 428–437.
[84] D. O'Sullivan, K. Ning, Semantic innovation management for the extended enterprise, Journal of Industrial and Systems Engineering 1 (1) (2006).
[85] A.H. Badradabi, S. Mohammadi, An assistant for large scale software project management: hybrid of agents and semantic web services, in: Proceedings of the IEEE International Conference on Management of Innovation and Technology, Bangkok, Thailand, 2008.
[86] S. Mobammadi, A. Khalili, A semantic web service-oriented model for project management, in: Proceedings of the 8th IEEE International Conference on Computer and Information Technology, Sydney, Australia, 2008.
[87] M. Tosic, V. Milicevic, M. Stankovic, Collaborative knowledge acquisition for agile project management, in: Proceedings of the International Conference on the Computer as a Tool (EUROCON 2005), 2005.
[88] H.K. Lin, J.A. Harding, M. Shahbaz, Manufacturing system engineering ontology for semantic interoperability across extended project teams, International Journal of Production Research 42 (24) (2004) 5099–5118.
[89] M. Chen, D. Zhang, L. Zhou, Empowering collaborative commerce with web services enabled business process management systems, Decision Support Systems 43 (2) (2007) 530–546.
[90] D. Fensel, Y. Ding, B. Omelayenko, E. Schulten, G. Botquin, M. Brown, A. Flett, Product data integration in B2B e-commerce, Intelligent Systems 16 (4) (2001) 54–59.
[91] W. Kim, M.J. Chung, J. Lloyd, Automated outsourcing partnership management using semantic web services, Computer Supported Cooperative Work in Design II (2006) 184–193.
[92] L. Li, I. Horrocks, A software framework for matchmaking based on semantic web technology, International Journal of Electronic Commerce 8 (4) (2004) 39–60.
[93] D. Moitra, J. Ganesh, Web services and flexible business processes: towards the adaptive enterprise, Information and Management 42 (7) (2005) 921–933.
[94] D. Shen, G. Yu, Y. Kou, T. Nie, Z. Zhao, Resolving heterogeneity of web-service composition in network manufacturing based on ontology, International Journal of Computer Integrated Manufacturing 20 (2) (2007) 222–233.
[95] J. Li, X. Meng, S. Liu, R. Yuan, Ontology-based resource description in manufacturing grid, in: Proceedings of the 11th International Conference on Computer Supported Cooperative Work in Design (CSCWD 2007), 2007.
[96] F. Liu, C. Ming, D. Jingxiang, Swirrsm: a manufacturing resource discovery and ranking system, in: Proceedings of the 11th International Conference on Computer Supported Cooperative Work in Design (CSCWD 2007), 2007.
[97] E.I. Neaga, J.A. Harding, H.K. Lin, Towards a meaningful manufacturing enterprise metamodel: a semantic driven framework, Proceedings of the Institution of Mechanical Engineers, Part B: Engineering Manufacture 221 (3) (2007) 499–511.
[98] S. Lemaignan, A. Siadat, J.Y. Dantan, A. Semenenko, MASON: a proposal for an ontology of manufacturing domain, in: Proceedings of the IEEE Workshop on Distributed Intelligent Systems: Collective Intelligence and its Applications (DIS'06), 2006.
[99] F. Colace, N. Pasquino, P. Senatore, Ontology for e-business: a semantic multi-level approach, in: Proceedings of the 3rd International Conference on Information and Communication Technologies, Damascus, Syria, 2008.
[100] Q. Thanh Tho, H. Siu Cheung, A.C.M. Fong, Automatic fuzzy ontology generation for semantic helpdesk support, IEEE Transactions on Industrial Informatics 2 (3) (2006) 155–164.
[101] A. Dimache, L. Dimache, E. Zoldi, T. Roche, Development of a strategy tool for environmental compliance management, in: Proceedings of the 17th CIRP Design Conference, Berlin, Germany, 2007.
[102] T. Roche, The design for environmental compliance workbench tool, in: Proceedings of the Advanced Summer Institute on Product Engineering, Eco-Design Technologies and Green Energy Conference, Brasov, Romania, 2004.
[103] M. Davis, Semantic Wave 2008 Report: Industry Roadmap to Web 3.0 and Multibillion Dollar Market Opportunities, http://www.researchandmarkets.com/reports/682250.

John Breslin is a lecturer in electrical and electronic engineering at the School of Engineering and Informatics, National University of Ireland, Galway. He is also an associate researcher and leader of the Social Software Unit at the Digital Enterprise Research Institute at NUI Galway, researching the Social Semantic Web. He is the founder of the SIOC project, which aims to semantically-interlink online communities. He has received a number of awards for website design, including two Net Visionary awards from the Irish Internet Association for the Irish community website boards.ie, which he co-founded in 2000. He is a member of the IEI, IET and IEEE.

David O'Sullivan is a lecturer and vice dean of teaching at the School of Engineering and Informatics, National University of Ireland, Galway. He is also a project leader at the Digital Enterprise Research Institute. His research interests are in the area of semantic innovation management and industrial applications of computer science. His recent projects have included semantic innovation across the extended enterprise and e-enablement of small- to medium-sized enterprises. His current work includes the use of Semantic Web technology for collaborative design environments.

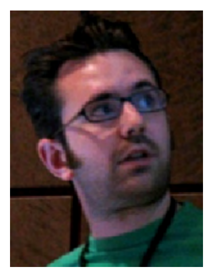

Alexandre Passant is a postdoctoral researcher at the Digital Enterprise Research Institute, National University of Ireland, Galway. His research activities focus around the Semantic Web and Social Software: in particular, how these fields can interact with and benefit from each other in order to provide a socially enabled machine-readable Web, leading to new services and paradigms for end-users. Prior to joining NUI Galway, he was a PhD student at Université Paris-Sorbonne and carried out applied research work on "Semantic Web technologies for Enterprise 2.0" at Électricité De France. He is the co-author of SIOC, a model to represent the activities of online communities on the Semantic Web, the author of MOAT, a framework to let people tag their content using Semantic Web technologies, and is also involved in various related applications as well as standardisation activities.

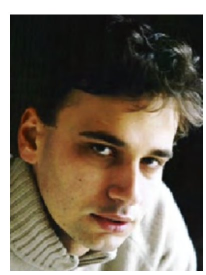

Laurentiu Vasiliu is a postdoctoral researcher and group leader with the Digital Enterprise Research Institute at the National University of Ireland, Galway. He has a PhD in industrial engineering and has been involved in several EU Semantic Web research projects including SWING, DIP and Tripcom. His interests are in applied research and commercialisation as well as B2B technologies, e-business and business intelligence.