

On the Emergent Semantic Web and Overlooked Issues

Yannis Kalfoglou¹, Harith Alani¹, Marco Schorlemmer², and Chris Walton³

¹ Advanced Knowledge Technologies (AKT), School of Electronics and Computer Science, University of Southampton, UK
{y.kalfoglou, ha}@ecs.soton.ac.uk

² Escola de Tecnologies d'Informació i Comunicació, Universitat Internacional de Catalunya, Spain
marco@cir.unica.edu

³ Advanced Knowledge Technologies (AKT), CISA, School of Informatics, University of Edinburgh, UK
cwd@inf.ed.ac.uk

Abstract. The emergent Semantic Web, despite being in its infancy, has already received a lot of attention from academia and industry. This resulted in an abundance of prototype systems and discussion most of which are centred around the underlying infrastructure. However, when we critically review the work done to date we realise that there is little discussion with respect to the vision of the Semantic Web. In particular, there is an observed dearth of discussion on how to deliver knowledge sharing in an environment such as the Semantic Web in effective and efficient manners. There are a lot of overlooked issues, associated with agents and trust to hidden assumptions made with respect to knowledge representation and robust reasoning in a distributed environment. These issues could potentially hinder further development if not considered at the early stages of designing Semantic Web systems. In this perspectives' paper, we aim to help engineers and practitioners of the Semantic Web by raising awareness of these issues.

1 Introduction

The *Semantic Web* (SW) is on its early stages of research and development – not more than five years in life loosely counting from appearances of early publications in the literature and W3C blueprints and design notes – and there has been considerable progress in both academic and industrial fronts. The emergent SW, however, still falls short of expectations and hasn't realised the vision outlined in [5]. This statement is bound to be subjective as many would interpret the “SW expectations” differently depending on their experiences with SW technologies and their successes or failures. In this paper we aim to put in perspective these views and critically review the SW vision and work done to date. We reflect on the expectations created by the SW literature, but our aim is not to conclusively discount the idea of the SW. Rather, we aim to raise awareness of overlooked issues with respect to: (a) provision of robust reasoning on the SW, (b) usage of mechanised forms of trust, (c) enabling software agents on the SW, and (d) cultural legacy from a Web user's point of view. We base our critique on personal experiences with involvement in the design, development, and deployment of SW-enabled applications, thus our discussion is targeted to practitioners and engineers and the focus is on technical issues and limitations of current technology.

From an ethnographical point of view, this sort of critique and analyses of a new technology is sparse in the SW literature. For example, in the past three SW events – starting with the SWWS'01 symposium and followed by two ISWC conferences – there

have been 135 papers published but, to our judgment, we only found 3 critiquing papers highlighting problems, experiences and issues emerged. Similar dearth of critiquing is found in SW related events organised by adjacent communities like the WWW or major Artificial Intelligence (AI) conferences. Based on this evidence and seeing all this excitement and work done so far, an average Web user would wonder: *the SW should be already out there and I could use it as I use the Web?*. However, this is not true yet, as current users of the emergent SW are mostly its engineers and practitioners rather than the general public.

These two are the main anchor points upon which we unravel our argumentation: review work done to date and highlight issues that could hinder future development by focusing on the expectations of the SW from a Web user's point of view. We start with a brief overview of the emergent SW in section 2 before elaborating on emerging issues in the following sections. In particular, in section 3 we argue for the role of agency and mechanised trust as the means to realise the benefits of a SW. These are closely related with our ability to operationalise robust reasoners on the SW, an issue related with Knowledge Representation (KR). We elaborate on known limitations of KR for the SW in section 4, and we put together all these issues under a pragmatics' lens in section 5 where we speculate on the impact of these issues and pinpoint to directions that could help us come closer to realise the SW vision before concluding our perspectives' paper in section 6.

2 The Emergent Semantic Web

The SW endeavour is well underway with the majority of work concentrating on infrastructure issues. As semantics are seen as the differentiating factor from the Web, one of the first deliverables of the SW, namely RDF [29], was aiming to assist content providers with annotation of their data with semantic information. The adoption and use of RDF though has not been unproblematic, as reported in [24]. In [42] the authors elaborate further on the problems associated with RDF syntax and notation, in particular when it is used with the W3C's recommended Web Ontology Language (OWL) [36]. OWL emerged in the last few years as an ontology language that, hopefully, will become a *de facto* language in which to represent SW ontologies. The design of this carefully crafted language was influenced by AI products, such as Frame Logic and Description Logic (DL). These formalisms have their own peculiarities which are inherited in the language itself but we will elaborate on them in the next section.

As RDF preceded OWL, early adopters of this technology used sophisticated, but ad-hoc ways of providing AI-style inferencing in early systems (for example, the *OntoBroker* system [10]). As OWL became widely available and backed by the W3C we are beginning to see content annotated with OWL syntax not only for ontology encoding but for describing instances too (ABoxes in DL). A number of tools are already available, including robust APIs like HP's Jena⁴ and the OWL-API⁵.

But despite the sophisticated machinery for annotating content and reasoning over it, the main driver of the SW seems to be simplicity and ease of use. This is evident

⁴ <http://jena.sourceforge.net>

⁵ <http://sourceforge.net/projects/owlapi>

since the early days of designing the SW [4], when the widespread belief was to keep things simple and to compromise on KR principles that shaped most of the KR practice in AI for years, as we will discuss in the next section. To quote Berners-Lee [4]: “The Semantic Web is what we will get if we perform the same globalisation process to Knowledge Representation that the Web initially did to Hypertext. We remove the centralised concepts of absolute truth, total knowledge, and total provability, and see what we can do with limited knowledge.” Similar statements are found in the SW literature almost 5 years after these design notes were written, as for example in McBride’s report [34]:

[...] the perception that the Semantic Web is concerned with AI is not helpful to its widespread adoption in the IT industry [...] this perception is understandable. The Semantic Web is often presented as the technology that will achieve marvelous things [...] much of the excitement and motivational power of such scenarios come from the handling of speech, NLP, general problem solving, scheduling, common sense reasoning and other features commonly associated with intelligence. These are not what the Semantic Web is about. The Semantic Web is about creating an infrastructure in which information from a variety of sources can be integrated on demand to achieve some task. The Semantic Web provides mechanisms that enable access to and integration of data.

There is a hidden antithesis in these two, chronologically distant, but similar calls for compromised KR and infrastructure. McBride’s call for infrastructure reflects much of the work done to date, but the infrastructure call presupposes a non-compromised KR, as the tasks envisaged for this infrastructure require the full strength of today’s best KR practices. You can’t have a compromised KR to support an infrastructure that aims to provide the means for accessing and integrating disparate data in semantically rich manners. At least not with the current state-of-the-art in KR and SW infrastructure. So, there is a compromise on both ends, KR and infrastructure, as it is evident on early, successful, SW systems and demonstrators.

For example, the 2003 SW challenge award-winner CS AKTive space⁶ strikes a balance between aggregation of vast amounts of semantically annotated data from a variety of disparate sources and provision of robust reasoning services. For instance, as the focus is on infrastructure, more attention has been paid on performance and scaling issues⁷ rather than robust reasoning, as for example resolving the ever so re-occurring problem of referential integrity [1]. On the other hand, systems like KAON [51], provide more robust reasoning services but compromise on scale and, arguably, on the principle of decentralization and distributiveness of the SW as KAON is a stand-alone application server with its processing centralised. The list of early prototypes who compromise on either scale and distributiveness or KR capabilities is long to mention in this paper but to the best of our knowledge none of those developed so far, have managed both.

An issue closely related with the ability to provide robust reasoning over distributed resources is not so much the capabilities of the reasoner but the usefulness of the exter-

⁶ <http://www-agki.tzi.de/swc/csaktivespace.html>

⁷ One of the biggest RDF triplestores available today with more than 25 million RDF triples in store.

nal resources. In an environment the size of the SW these could be assessed by some form of mechanised trust and accessed with software agents, as it is envisaged in the literature. There are open issues, however, when we call upon agency and trust which we review in the next section.

3 The Role of Agency and Trust

The SW is predicated on the notion of *agents* as the key consumers of SW information. This is expressed in [5]: “The real power of the Semantic Web will be realised when people create many programs that collect Web content from diverse sources, process the information and exchange the results with other programs. The effectiveness of such software agents will increase exponentially as more machine-readable Web content and automated services (including other agents) become available.”. The underlying assumption in this description is that agents will be able to *automatically* utilise the information on the Web once semantic annotation has been provided. Nonetheless, there are many challenges associated with the use of agency in this manner, and a solution to many of the issues is far from clear. We outline a number of these challenges in this section, with reference to state-of-the-art techniques for addressing the associated issues.

What is an agent? The first issue concerns the definition of the term “software agent” itself. There is a surprising lack of consensus in the agency community over how this term should be defined. Indeed, the term is applied to a wide range of systems, ranging from simple distributed objects, to complex autonomous entities [41]. The reason for this lack of consensus is primarily because the term is used generically to refer to a heterogeneous body of AI research, rather than a specific implementation. From the perspective of the SW, it is unclear what kind of agency is required. The increasing use of Web services⁸ to express computation on the SW points to a purely procedural notion of agency, while the kinds of reasoning which are envisaged in the description of the SW appear to require something more complex, e.g. proactive behaviour. Thus, to address this issue, a number of systems that combine autonomous agents with Web services have been proposed (e.g., [52]).

Agent Coordination. It can be argued that a precise definition of agency on the SW is unimportant, provided that the agents can cooperate in meaningful ways. This is further expressed in [5]: “[...] even agents that were not expressly designed to work together can transfer data among themselves when the data come with semantics.”. From this description it appears that the semantic markup itself is intended to solve the problem of coordination among heterogeneous agents. However, it is far from clear how this can be achieved. The problem of coordinating agents with semantic markup introduces the issues associated with semantic interoperability (see section 4). For example, if two agents wish to cooperate, and they have different internal representations of the domain in question, then they must perform a mapping between (or a subset of) their knowledge. This must be done automatically if the agents are to cooperate autonomously in the

⁸ <http://www.w3.org/TR/ws-arch/>

SW. There are a number of proposed solutions, e.g., the adoption of a standard upper ontology such as OWL-S⁹, or the use of partial mappings [23]. However, it is likely that this will be an open issue for some time to come.

A popular basis for coordination among agents is the use of *performative* languages. These languages have their roots in the theory of *speech acts*. This theory identifies a class of performative verbs (e.g. inform, promise, request) that have the characteristics of physical actions, in that they change the state of the world. The use of performatives to express interactions between agents has been enthusiastically adopted by the agency community and this is most visible in the development of Agent Communication Languages (ACLs), such as the Foundation for Intelligent Physical Agents ACL (FIPA-ACL¹⁰). In these languages, the model of interaction between agents is based on the exchange of messages. FIPA-ACL defines a set of performatives (message types) that express the intended meaning of these messages.

Sincerity. There is a significant issue in the use of performative languages in an open environment. There is an underlying *sincerity assumption* in these languages which demands that agents always act in accordance with their intentions. It will always be possible for an insincere agent to simulate any required internal state, and we cannot verify the sincerity of an agent as we have no access to its internal mental states. This problem is termed the *semantic verification* problem and is detailed in [54]. In order to address the sincerity issues, a number of alternatives have been proposed which we now discuss. The proposals all address sincerity by restricting the kinds of behaviours that the agents can exhibit and exposing the internal state of the agents. Thus, agents who are not acting in accordance with the intended dialogue can be readily identified.

Dialogue protocols. The first of these imposes a layer of control above the performative language, through the use of *dialogue protocols* (e.g., [49]). In this approach, the performative language defines only *how* agents can communicate, while the dialogue protocol defines *if and when* agents should communicate. The underlying concept in the Electronic Institutions (EI) approach is that human interactions are never completely unconstrained, rather they are always guided by formal and informal conventions such as customs, etiquette, and laws. EI are a means of representing and controlling these conventions within multi-agent systems (MAS). Dialogue protocols greatly assist in the design of MAS as they impose structure on the agents, co-ordinate tasks between agents, and define commitments which agents must satisfy. Nonetheless, there are a number of pertinent issues which need to be addressed in the EI framework. The most serious of these is that there is no straightforward means to disseminate the institutions, and thus the agents must be explicitly designed (in advance) for a specific institution.

Social policy. another alternative is the use of social policy to control agent interaction. The key concept of the social commitment model is the establishment of shared commitments between agents. A social commitment between agents is a binding agreement from one agent to another. The commitment distinguishes between the creditor who commits to a course of action, and the debtor on whose behalf the action is done.

⁹ <http://www.daml.org/services/>

¹⁰ <http://www.fipa.org>

Establishing a commitment constrains the subsequent actions of the agent until the commitment is discharged. Commitments are stored as part of the social state of the MAS and are verifiable. A theory which combines speech acts with social commitments is outlined in [17], though there remain considerable implementation issues in maintaining a suitable store of shared commitments.

Dialogue games. a further approach involves the use of dialogue games, which trace their origins to the philosophical tradition of Aristotle. Dialogue games have been used to study fallacious reasoning, for natural language processing and generation, and to develop a game-theoretic semantics for various logics. These games can also be utilised as the basis for interaction between autonomous agents. A group of agents participate in a dialogue game in which their utterances correspond to moves in this game. Different rules can be applied to the game, which correspond to different dialogue types, e.g. persuasion, negotiation, enquiry [53]. For example, a persuasion dialogue begins with an assertion and ends when the proponent withdraws the claim or the opponent concedes the claim. A framework which permits different kinds of dialogue games, and also meta-dialogues is outlined in [35]. However, this framework is only defined at a formal level, and an implementation has yet to be realised.

Dynamic scripting. the final approach which we describe here relaxes the requirement that the agent coordination must be statically defined before evaluation can take place. Rather, we can take a script-based approach to coordination, where the script is built dynamically during evaluation, as the agents communicate. In general, a great deal of conversation is an expression of protocol. For example, the question “*Can you tell me the time?*” implies that the next step in the protocol is for the other party to respond with the time. We express the coordination as an executable protocol, which is passed between agents as the communication happens. This approach is detailed in [47], where the scripting language is derived from the Calculus of Communication Systems [39].

Trust

One of the needs for software agents’ cooperation on the SW is the ability to use automated methods for measuring trust. This is exemplified in [5]: “Lucy’s agent, having complete trust in Pete’s agent in the context of the present task, automatically assisted by supplying access certificates...” *Trust* is placed at the top of the proposed SW “layer cake”¹¹. The trustworthiness of knowledge-base statements have generally been accepted without any proof [50]. Van Harmelen argues that some measurement of trust will be needed when dealing with distributed knowledge sources. In [32] the authors highlight some of the issues that need to be incorporated into the SW service architecture to enable exchanging trust and reputation and to control policies and negotiations. There are several issues to consider with respect to trust and the SW. For example, *how can trust be modelled and exchanged between agents and SW services? Where should trust annotations be stored and made available? What kind of knowledge is required to measure trust and where will this knowledge come from? What trust features need to be considered (e.g. subjectivity, propagations, transitivity)? and how do they effect trust in general?*

¹¹ <http://www.w3c.org/2000/Talks/1206-xml2k-tbl/slide10-0.html>

Modelling & representation. For trust to be used by SW agents it must first be formalised and modelled in an agreed fashion. The first computational model of trust was introduced in [31], which inspired much work in this domain. Modelling trust helps to untangle its complexity and understand its main components and evolution. However, the accuracy and completeness of such models is one thing, and their suitability and integration into MAS communication is a totally different thing. In [19] the authors are calling for security and trust to be added to the systems design process, rather than trying to attach them to already existing systems which might not be compatible. They argue that modelling trust and security should be closely related to where they will be applied and used, which raises questions about the actual applicability of generic trust models.

There is a clear need for semantic language support for representing trust. In [16] the authors suggested to use a semantic policy language to enable annotating security and belief. Representing trust explicitly can significantly improve knowledge exchange and outsourcing performance between agents [13]. There are some emerging standards such as SOAP security extension and PICS which was originally designed for rating Web pages. However, none of these languages are sophisticated enough to represent trust efficiently.

Sourcing. There is still no agreement about whether trust information is best gathered and broadcasted in centralised units, or left in a distributed format for agents to collect and reach their own conclusions. In [18] the authors developed a centralised system to help users annotate sources to highlight alternative, probably contradictory, sources. Users feedback can then be used to assess information sources. Another example is proposed in [33] for a centralised agent to measure the reputation of Web services by monitoring and collecting outcomes and client feedback, and making this information available to other agents. Such centralised systems raise the question of how trustworthy are the users or agents who provided the feedback in the first place, and why such trust warehouses should be trusted [13]. For example, in [16] the authors argue against centralised units for measuring trust because of their scalability limitations and the implicit trust measurement mechanisms they adopt.

Measuring trust. to reduce the risk of online trading on the SW, there must be a way to infer trust and reputation electronically [40]. However, there is no consensus about how this can be done. In [55] the authors state the need for agents to combine their own experience with the experience of other agents when measuring the trustworthiness of a specific agent. This raises the issue of availability and accessibility of such experience outcome over the SW, and whether trust can still be measured in their absence.

Several approaches have been introduced in an attempt to extract trust values from decentralised, interconnected semantic networks. In [22] the authors argue that trust can be propagated through relatively short paths connecting entities in small world networks. They believe that modern information networks will be decentralised and highly dynamic, where traditional trust control methods will no longer apply. They used an algorithm based on the *small world theory* [38] to measure trust between entities that are not directly connected. Similarly, in [20] the authors applied techniques to measure trust over a FOAF network, extended with trust relations while in [46] the authors applied path algebra and probabilistic algorithms to measure the trustworthiness of users

along a network of trust connections explicitly specified by users. All these approaches seem to rely to some extent on trust values given explicitly by users. One observation made in [45] with respect to *eBay* is that users feedback is almost always positive. The authors note that most people do not like giving negative feedback, unless revenge is a motivation. Furthermore, scalability becomes a serious issue when relying solely on users trust feedback.

Context. one major problem with many approaches proposed for measuring trust is the lack of context, which is a basic feature of trust [13]. You may trust your colleague to write a good project proposal, but you might not trust him to fix your car. Associating trust measurement with specific contexts will inevitably increase its calculation and representation complexity, but nevertheless is crucial for the resulting trust values to be on any use. Transitivity of trust is a complex issue that also needs to be investigated to determine how trust should propagate along the SW. Some argue that trust is not strictly transitive, and that the transitivity of trust is subjective and context dependent [12].

It should now be clear that the capitalisation of agency and trust in the SW will require a significant number of challenges to be addressed. The focus of our attention has been on the provision of a coordination framework for agents, and mechanised ways for measuring trust. However, even with such a framework in place and a well-connected network upon which network analysis algorithms will measure trust, we are still dependent on robust reasoners that will operate in such an environment.

4 Knowledge Representation and Reasoning on the Semantic Web

Robust reasoning. Intimately connected with the ability to operationalise our reasoners on the SW is how we handle *soundness* and *completeness*. These are precise technical terms that describe properties of formal systems or sets of sentences, and they get their meaning from the field of mathematical logic. A formal system is *sound* when every sentence produced by the system's inference rules operating on the system's initial set of axioms logically follows from that set; it is *complete* when every sentence that logically follows from the system's initial set of axioms can be formally derived using the inference rules. In addition a set of sentences is said to be *complete* if for every sentence of the logical language, either it or its negation can be proven from the set, using the rules of the formal system.

These notions are tightly connected to the story of first-order logic (FOL), which, among all logics, has a special status due to its expressive power, its natural deductive systems, and its intuitive model theory based on sets. It is because of FOL's success in providing rigorous definitions of mathematical truth and proof, that, from the early days of the field of AI, knowledge representation and reasoning (KRR) has been developed primarily around FOL and its model-theoretic semantics.

From a practical point of view, to render the SW operational, it is understandable that current standardisation efforts for SW technology have primarily focused around traditional views of KRR, using FOL's model theory as its semantic underpinning. FOL is well understood, both mathematically and computationally, and is also well established via many customised representation languages and highly optimised inference

engines that have been developed for a plethora of different subsets and variants of FOL. Hence, it comes at no surprise that OWL is based on well-studied DL theory and technology, very well following FOL's traditional model-theoretic semantics [36].

But, despite of FOL's prominent status in mathematical logic, and besides the practical reasons for adopting FOL's model theory as semantic underpinning to get the SW up and running, it would be a mistake to base the current emergent SW technology solely, or primarily, around the traditional KRR paradigm. In [15] the author has nicely summarised the strengths and limitations of the model-theoretic approach to the SW, which situates many of the envisioned applications of the SW clearly beyond the capacity of FOL-based, and consequently DL-based SW technology.

Soundness and completeness, for instance, have a different dimension when it comes to the SW. As it is reported in [50]: "[...] some of the assumptions underlying current KR technology [that] will have to be revised when applied to the Semantic Web." While scale and change rate are regarded as big obstacles that test current KR technology to its limits, the crux of the problem with preserving and mechanizing *soundness* and *completeness* on the SW is lack of referential integrity and inconsistent knowledge produced by multiple knowledge sources, as they will exist in the SW. There is call for more robust inferencing [50]:

In an environment the of size of the web we must abandon the classical idea of sound and complete reasoners, our reasoners will almost certainly have to be incomplete (no longer guaranteeing to return all logically valid results), but most likely also unsound: sometimes jumping to a logically unwarranted conclusion. Furthermore, the degrees of such incompleteness or unsoundness must be a function of the available resources. Answers will have to be approximate.

Open and closed worlds. We visit a key notion that has been used in the past to tackle incompleteness, and it is revised in the context of agents' research to become more operational when applied at a large scale. *Closed World Assumption* (CWA) states that everything that is not known or cannot be proved to be true, must be false. CWA originates from AI and database research in the late 70s (for example, Reiter's work on CWA [44] as an implementation of McCarthy's investigations on the *Frame Problem*) and the same fundamental assumption still holds for most of databases' design today. Despite its similarities, it has a different interpretation when studied in DL, as it is stated in [2]: "a database instance represents exactly one interpretation, namely the one where classes and relations in the schema are interpreted by the objects and tuples in the instance [...] an ABox represents many different interpretations, namely all its models." So, absence of information in a database instance is interpreted as *negative information*, while absence of information in an ABox only indicates *lack of knowledge*. Thus, while the information in a database is always understood to be complete, the information in an ABox is in general viewed as incomplete, and that leads to the perception of ABoxes semantics as "open world", while the traditional semantics of databases is characterised as "closed world".

This view that ABoxes (set of instances or facts that comprise a knowledge base) have an "open world" semantics is clearly reflected in the current SW technologies which originate from DL (for instance, in OWL). For this reason, CWA has received a

rather controversial perception. In the W3C OWL requirements recommendation [25], the ability to state closed worlds is seen as an objective, rather than a requirement¹² and its value is praised by the author: “[...] the language must be able to state that a given ontology can be regarded as complete. This would sanction additional inferences to be drawn from that ontology. The precise semantics of such statement (and the corresponding set of inferences) remains to be defined, but examples might include assuming complete property information about individuals, assuming completeness of class-membership, and assuming exhaustiveness of subclasses.”. These features are indeed what a SW reasoner would expect to find in a knowledge source on the SW in order to draw complete, and possibly sound inferences.

But, providing the right machinery to automate the task of asserting information about the completeness of a knowledge source is not an easy job and solutions vary across disciplines. For example, in the logic programming realm, the use of extra-logical operators, embedded in the syntax of the language, are common, like the use of *not* operator in some implementations of Prolog. Similar attempts are reported in the software engineering domain [30], whereas in [6] the authors use techniques from Answer Set Programming (ASP), which allow to draw conclusions based on the lack of evidence of the contrary. Their approach augments the semantics of SW languages’ constructs, in particular that of the DAML `daml:subClassOf`, with ASP programs to be fed into an ASP solver in order to allow drawn, default inferences to be dropped if opposite knowledge is found in the knowledge source. A promising approach emerged in the agents’ community as an extension of the classical CWA, namely *Local World Assumption* - hereafter, LCW. It was first applied in the XII planner [21] to reduce the planning time when dealing with external information. LCW allows closed-world information to be obtained on subsets of information that are known to be complete, while still allowing other information to be treated as unknown. A problem acknowledged with early LCW approaches though, was the assumption that there must exist a priori knowledge of the local completeness of the information sources. And that assumption cannot be operational in an environment the size and scale of the Web. In [26] a proposed solution to this is to provide the means for facilitating the assertion of local completeness information in SW languages themselves, thus enabling SW content providers to annotate their resources with LCW statements. However, the issue of how to acquire knowledge of local completeness in the first place still remains unsolved. The plausible assumption made here, is that the content providers should possess that knowledge and they only need to annotate their content with it. Still, this leave some open issues with respect to the consistency of their content, as LCW could easily introduce contradictions if not handled properly when applied at huge resources.

Semantic interoperability. Spurred by the need to be able to reuse knowledge-based systems and to make them interoperable, the effort in achieving knowledge sharing has carried this tradition in KR over to the research on information system ontologies. In such a highly distributed and decentralised environment as that envisioned for the SW, ontologies were quickly adopted as its semantic backbone: The proposed layer di-

¹² The difference being, that a W3C requirement is a feature that the OWL language must have as opposed to an objective which is not mandatory but nice to have.

agram¹³ reflecting the SW architecture includes a layer for “ontological vocabulary”. But the proliferation of ontologies and ontology representation languages has highlighted the need of standardisation, similarly to that achieved at lower levels of the proposed architecture diagram, which encapsulate already *de facto* standards of current Web technology, such as URI’s and XML syntax. As already mentioned in section 2, in the past years RDF and RDF Schema, with their model-theoretic semantics, have emerged as recommendations for formalisms for defining the semantic content of the Web [29].

Within ontologies’ field, many proposals for standard ontologies have arisen in various domains, and the difficulty of consensus has not hindered the continuing work on even the old dream of philosophy of building representations of common-sense knowledge in so called “upper” or “top-level” ontologies such as Cyc¹⁴ and IEEE’s Standard Upper Ontology.¹⁵ These efforts attempt to provide generic “objective” representations of knowledge about our world on which ideally particular domain-specific knowledge would define their primitive terminology.

Without questioning the value and necessity of standardisation efforts of this kind, such centralised ontologies may basically make sense for small communities. But large standard ontologies bring back the rigidity of centralised agencies organised under classical codification-based management philosophies [9]. The very distributed and decentralised nature of the web makes the proliferation of ontologies and of formalisms to represent these an unavoidable reality. This has put forward the need to support the interoperability and integration of communities on the semantic level, for instance by means of technology that supports the mapping and merging of ontologies [28].

Corrêa da Silva et al. have shown situations in which even a shared conceptualisation as explicitly specified by means of an ontology is not enough for achieving flawless knowledge sharing [8]: A knowledge base whose inference engine is based on linear logic that poses a query to a knowledge base with the same ontology, but whose inference engine is based on relevance logic should not accept answers as valid if the inference carried out in order to answer the query was using the contraction inference rule, which is not allowed in linear logic. Here, we have clearly an issue of *trust* on the other knowledge base’s inference engine, which results from each knowledge base being based on different logical systems.

In [15] Farrugia suggests that, before any meaning negotiation between two agents can start, first the “logical setup” needs to be established, for which he points to Meseguer’s notion of *logical system* [37]. This contrasts with the mainstream effort within the SW community where semantic interoperability is sought within the context of standardised representation languages based on Tarskian model-theoretic semantics.

Alternative approaches for a logic based on precise mathematical models of information as a necessary requirement for designing and operating information-processing systems have been advocated e.g., [11, 3]. We have recently explored how mathematical theories of information may provide a different angle from which to approach the distributive nature of semantics on the Semantic Web. As it seems, an information-theoretic

¹³ <http://www.w3c.org/2000/Talks/1206-xml2k-tbl/slide10-0.html>

¹⁴ <http://www.cyc.com>

¹⁵ <http://suo.ieee.org>

approach such as that of Barwise and Seligman's channel theory [3] may be suitable to accommodate various understandings of semantics like those occurring in the Web, and also allows for the establishing of ontology mappings that accommodate the particular way different communities use their terminology in their respective contexts, as defined by the use and classification of particular instances with respect to local concepts [27].

5 Pragmatics

In sections 3 and 4 we reviewed and analysed a number of challenges that SW faces today. In this section we focus on their impact and put them in perspective when we consider short to medium term deliverables. As the challenges we reviewed cover a broad area of scientific research, it is not realistic to expect them to be fully resolved before the SW will be available and commercially exploitable. It will take time to come up with sound scientific and practical solutions to the problems of robust reasoning, agency coordination, and semantic interoperability, to name a few. In the meantime, the SW will continue to grow and attract attention based on short to medium term solutions. This model of evolution follows the Web legacy, when the SW's predecessor grew on the basis of a network effect. Early Web adopters volunteered to link up their content which resulted in the biggest distributed network ever build in modern computer science. This fundamental assumption is passed onto SW adopters as linked content is still the main deliverable. The difference which should be catered for, though, is that a broken semantic link in a futuristic SW will take more than an explanatory Web-style "404" error message to resolve. As mechanised reasoners will be the main crawlers of content, and not humans as it is often the case with the Web, there have to be automated ways of getting around broken semantic links. Safeguarding your inferences against such an event will be crucial for the quality of service a SW provider will aim to deliver.

Another consideration with respect to the network effect principle of building the SW, is that content management is no longer a carefully manifested process controlled by few but rather an ad-hoc procedure engaged by many. The side-effect of this, is that trust and interoperability will take centre stage as the means to ensure that only useful content will be fed into our reasoners. We also see from evidence from analysing the behaviour of users of the Internet as whole, that erosion of trust is something that we should become to expect in the future, as Blumenthal and Clark point out: "Of all the changes that are transforming the Internet, the loss of trust may be the most fundamental. The simple model of the early Internet - a group of mutually trusting users attached to a transparent network - is gone for ever. A motto for tomorrow may well be 'global communication with local trust'." [7]. Providing the means for local trust would relax the "tension between the demand for trustworthy overall operation and the inability to trust the behaviour of individual users"[7]. As we reported in section 3, laying out the technological foundations for operationalising this 'local trust' that Blumenthal and Clark call for is a whole new challenge in its own right.

The community's favourable route to provide reliable reasoning, is to build as robust as possible reasoning machines for the SW. This is clearly reflected in the latest OWL family of tools (e.g., OWL languages, OWL rules, OWL-S) but would probably not be enough to tackle flawed content. As we reported in section 4 the current state-of-

the-art in KRR prevents us from providing guaranteed (i.e., complete) and trustworthy (i.e., sound) automated reasoning over incomplete and/or inconsistent content.

Furthermore, in [48] it is reported that the focus on the KR formalisms RDF and OWL does not include current computer technology, such as the query and modelling languages SQL and UML, into the SW in a satisfactory way. Such languages are widely used in database and software engineering circles, and they are more expressive than OWL. Any pragmatic approach to the SW would also have to aim at a seamless integration of this technology into the SW architecture.

Lastly, but not least, we have to consider cultural issues related to users' expectations and *modus operandi*. As the SW is seen by many as the natural evolution of the Web, users' will become to expect the same, and probably, more advanced usability and functionality. For example, inspired by the unprecedented success of the Web browser, some have engaged in exploratory activities to build "semantic web browsers" (e.g., [14, 43]). However, these are far from being fully operational or at least comparable with current Web browser technology. And we believe that this is a crucial point when it comes to promoting the SW idea to Web users as they have accumulated a 10 year experiences with the Web and only a truly superior product will win them over.

6 Conclusions

Despite using a pessimistic tone in pragmatics' section we are optimistic of the SW idea. We do, however, speculate on the process of building such an ambitious network and argue for a change of course. It seems more practical, and arguably easier to build and deliver in short time, to change the process of building *a* SW, to that of building a series of *Semantic Webs*, all of which will be smaller in size than the envisaged SW, but at least deliver – to the maximum degree possible – the vision outlined in Berners-Lee et al.'s Scientific American article. It could be argued that early SW prototypes are such small *Semantic Webs*. Once we master the art of managing these *Semantic Webs* and they reach a critical mass with satisfied users, the next natural step will be to link them, thus creating a *Web of Semantic Webs*, which eventually could evolve to the envisaged SW. This course is contrary to the current trend of hastily annotating with dubious semantics masses of content in the hope that advanced, but yet-to-be-built, reasoners and intelligent services will exploit this content and deliver the output to eager users, thus creating the SW. Which approach is the best, time can only tell but it's worth exploring both.

Acknowledgments. This work is supported under the Advanced Knowledge Technologies (AKT) Interdisciplinary Research Collaboration (IRC), which is sponsored by the UK Engineering and Physical Sciences Research Council under grant number GR/N15764/01. The AKT IRC comprises the Universities of Aberdeen, Edinburgh, Sheffield, Southampton and the Open University. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing official policies or endorsements, either express or implied, of the EPSRC or any other member of the AKT IRC. M. Schorlemmer is supported by a *Ramón y Cajal* Fellowship from the Spanish Ministry of Science and Technology.

References

1. H. Alani, S. Dasmahapatra, N. Gibbins, H. Glasser, S. Harris, Y. Kalfoglou, K. O'Hara, and N. Shadbolt. Managing reference: Ensuring Referential Integrity of Ontologies for the Semantic Web. In *Proc. of 13th Int. Conf. on Knowledge Engineering and Knowledge Management (EKAW'02)*, pages 317–334, 2002.

2. F. Baader, D. Calvanese, D. McGuinness, D. Nardi, and P. Pater-Schneider. *The Description Logic Handbook*. Cambridge University Press, 2003.
3. J. Barwise and J. Seligman. *Information Flow*. Cambridge University Press, 1997.
4. T. Berners Lee. What the Semantic Web can represent. *Web design Issues*, 1998.
5. T. Berners-Lee, J. Hendler, and O. Lassila. The Semantic Web. *Scientific American*, May 2001.
6. E. Bertino, A. Provetti, and F. Salvetti. Local Closed-World assumptions for reasoning about Semantic Web data. In *Proc. of APPIA-GULP-PRODE Joint Conference on Declarative Programming, AGP'03*, 2003.
7. M. Blumenthal and D. Clark. Rethinking the Design of the Internet: The end-to-end arguments vs. the brave new world. *ACM Transactions on Internet Technology (TOIT)*, 1(1):70–109, 2001.
8. F. Corrêa da Silva, W. Vasconcelos, D. Robertson, V. Brilhante, A. de Melo, M. Finger, and J. Agustí. On the insufficiency of ontologies: Problems in knowledge sharing and alternative solutions. *Knowledge Based Systems*, 15(3):147–167, 2002.
9. F. da Silva and J. Agustí. *Knowledge Coordination*. John Wiley and Sons Ltd., 2003.
10. S. Decker, M. Erdmann, D. Fensel, and R. Studer. Ontobroker: Ontology Based Access to Distributed and Semi-Structured Information. In *Proc. of DS-8, Semantic Issues in Multimedia Systems*, pages 351–369, 1999.
11. K. Devlin. *Logic and Information*. Cambridge University Press, 1991.
12. T. Dimitrakos and J. Bicarregui. Towards a Framework for Managing Trust in e-services. In *Proc. of 4th Int. Conf. on Electronic Commerce Research*, pages 360–381, 2001.
13. L. Ding, L. Zhou, and T. Finin. Trust based Knowledge Outsourcing for Semantic Web agents. In *Proc. of 2nd International Conference on Web Intelligence (WI-2003)*, 2003.
14. M. Dzbor, J. Domingue, and E. Motta. Magpie—Towards a Semantic Web browser. In *Proc. of 2nd International Conference on the Semantic Web (ISWC'03)*, pages 690–705, 2003.
15. J. Farrugia. Logical Systems: Towards protocols for Web-based meaning negotiation. In *Proc. of AAAI 2002 Workshop on Meaning Negotiation*, pages 56–59, 2002.
16. T. Finin and A. Joshi. Agents, Trust, and Information access on the Semantic Web. Technical report, Dec. 2002. SIGMOD Record, Volume 31, Number 4.
17. R. A. Flores and R. C. Kremer. Bringing Coherence to Agent Conversations. In *Proc. of Agent-Oriented Software Engineering (AOSE 2001)*, LNCS 2222, pages 50–67, 2002.
18. Y. Gil and V. Ratnakar. Trusting Information sources: One citizen at a time. In *Proc. of 1st Int. Conf. on the Semantic Web (ISWC'02)*, pages 162–177, 2002.
19. P. Giorgini, F. Massaci, J. Mylopoulos, and N. Zannone. Requirements engineering meets trust management: model, methodology, and reasoning. In *Proc. of 2nd Int. Conf. on Trust Management (iTrust'04)*, 2004.
20. J. Golbeck, B. Parsia, and J. Hendler. Trust networks on the Semantic Web. In *Proc. of 7th Int. Workshop on Cooperative Information Agents (CIA-2003)*, 2003.
21. K. Golden, O. Etzioni, and D. Weld. Omnipresence without omnipresence. In *Proc. of 12th Nat. Conf. on Artificial Intelligence (AAAI'94)*, 1994.
22. E. Gray, J. Seigneur, Y. Chen, and C. Jensen. Trust propagation in small world. In *Proc. of 1st Int. Conf. on Trust Management (iTrust'03)*, 2003.
23. M. Burstein. Ontology Mapping for Dynamic Service Invocation on the Semantic Web. In *Proc. of AAAI Spring Symposium on Semantic Web Services*, 2004.
24. S. Haustein and J. Pleumann. Is participation in the Semantic Web too difficult? In *Proc. of 1st Int. Conf. on the Semantic Web (ISWC'02)*, pages 448–454, 2002.
25. J. Heflin. OWL Web Ontology Language use cases and requirements. Technical report, W3C, 2004. W3C Recommendation—10 Feb 2004.
26. J. Heflin and H. Munoz-Avila. LCW-based agent planning for the Semantic Web. In *Proc. of AAAI'02 Workshop on Ontologies and the Semantic Web (WS-02-11)*, 2002.
27. Y. Kalfoglou and M. Schorlemmer. IF-MAP: an ontology mapping method based on Information Flow theory. *Journal on Data Semantics*, 1:98–127, LNCS 2800, 2003.
28. Y. Kalfoglou and M. Schorlemmer. Ontology mapping: the state of the art. *The Knowledge Engineering Review*, 18(1):1–31, 2003.

29. O. Lassila and R. Swick. Resource Description Framework(RDF) Model and Syntax Specification. W3C recommendation, W3C, Feb 1999.
30. Luqi and D. Cooke. How to combine nonmonotonic logic and rapid prototyping to help maintain software. *Int. J. of Software Engineering and Knowledge Engineering*, 5(1):89–118, 1995.
31. S. Marsh. Formalising trust as a computational concept. PhD thesis, University of Stirling, 1994.
32. B. Matthews and T. Dimitrakos. Deploying trust policies on the Semantic Web. In *Proc. of 2nd Int. Conf. on Trust Management (iTrust'04)*, 2004.
33. M. Maximilien and P. Singh. An ontology for Web service ratings and reputations. In *Proc. of AAMAS 2003 Workshop on Ontologies in Agent Systems (OAS'03)*, 2003.
34. B. McBride. Four steps towards the widespread adoption of the Semantic Web. In *Proc. of the 1st Int. Conf. on the Semantic Web (ISWC'02)*, pages 419–423, 2002.
35. P. McBurney and S. Parsons. Games that agents play: A formal framework for dialogues between autonomous agents. *Journal of Logic, Language and Information*, 11(3):315–334, 2002.
36. D. McGuinness and F. van Harmelen. OWL Web Ontology Language. Technical report, W3C, 2004.
37. J. Meseguer. Formal Interoperability. In *Proc. of 5th Int. Symp. on Artificial Intelligence and Mathematics*, 1998.
38. S. Milgram. The small world problem. *Psychology Today*, pages 60–67, 1967.
39. R. Milner. *Communication and Concurrency*. Prentice Hall International, 1989. ISBN: 0131150073.
40. L. Mui, M. Mohtashemi, and A. Halberstadt. A computational model for trust and reputation. In *Proc. of 35th Hawaii Int. Conf. on System Sciences (HICSS-35)*, 2002.
41. H. S. Nwana. Software Agents: An Overview. *The Knowledge Engineering Review*, 11(3):1–40, 1996.
42. P. Patel-Schneider and D. Fensel. Layering the Semantic Web: Problems and directions. In *Proc. of 1st Int. Conf. on the Semantic Web (ISWC'02)*, pages 16–30, 2002.
43. D. Quan and D. Karger. How to make a Semantic Web browser. In *Proc. of 13th Int. World Wide Web Conf. (WWW'04)*, 2004.
44. R. Reiter. On Closed World databases. In *Proc. of 1978 ACM SIGMOD Int. Conf. on Management of Data*, 1978.
45. P. Resnick and R. Zeckhauser. Trust among strangers in Internet transactions: Empirical analysis of eBay's reputation system. *Advances in Applied Microelectronics*, 11, 2002.
46. M. Richardson, R. Agrawal, and P. Domingos. Trust management for the Semantic Web. In *Proc. of 2nd Int. Conf. on the Semantic Web (ISWC'03)*, pages 351–368, 2003.
47. D. Robertson. A Lightweight Method for Coordination of Agent Oriented Web Services. In *Proc. of the AAAI Spring Symposium on Semantic Web Services*, 2004.
48. J. Sowa. Architectures for Intelligent Systems. *IBM Systems Journal*, 41(3), pages 331–349, 2002.
49. J. A. Rodríguez, F. J. Martín, P. Noriega, P. Garcia, and C. Sierra. Towards a Formal Specification of Complex Social Structures in Multi-Agent Systems. In *Proc. of Collaboration between Human and Artificial Societies conference*, pages 284–300, 2000.
50. F. van Harmelen. How the Semantic Web will change KR: challenges and opportunities for a new research agenda. *The Knowledge Engineering Review*, 17(1):93–96, 2002.
51. R. Volz, D. Oberle, S. Staab, and B. Motik. KAON server: A Semantic Web management system. In *Proc. of the 12th Int. World Wide Web Conf. (WWW'03)*, 2003.
52. C. Walton. Model Checking Multi-Agent Web Services. In *Proc. of the AAAI Spring Symposium on Semantic Web Services*, 2004.
53. D. N. Walton and E. C. W. Krabbe. *Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning*. SUNY Press, 1995.
54. M. Wooldridge. Semantic issues in the verification of agent communication languages. *Autonomous Agents and Multi-Agent Systems*, 3(1):9–31, 2000.
55. B. Yu and M. Singh. An evidential model of distributed reputation management. In *Proc. of the 1st Int. Joint Conf. On Autonomous Agents and Multi-Agent Systems (AAMAS'02)*, pages 294–301, 2002.