Semantic Social Software: Semantically Enabled Social Software or Socially Enabled Semantic Web?

Sebastian Schaffert

Salzburg Research Forschungsgesellschaft, Salzburg NewMediaLab
Jakob Haringer Str. 5/II, A-5020 Salzburg, Austria
sebastian.schaffert@salzburgresearch.at

Abstract. Semantic Social Software – the combination of social software and Semantic Web technology – has recently been gaining significant attention in the Semantic Web community. This article first discusses the two different perspectives the research community takes on "Semantic Social Software", namely the "Semantically Enabled Social Software" – meaning social software enhanced by semantics, and the "Socially Enabled Semantic Web" – meaning exploiting social software properties for creating Semantic Web data. It then introduces three example applications of Semantic Social Software (Semantic Wikis, Semantic Weblogs, and ePortfolios) and derives challenges for the Semantic Web community to address in order to overcome open issues of both, Semantic Social Software and the Semantic Web.

1 Introduction: Connecting Social Software and the Semantic Web

Recently, the combination of Social Software with Semantic Web technology has been gaining significant attention in the Semantic Web community. This is exemplified by the surprisingly high number of submissions and attendees of the *1st Workshop on Semantic Wikis* [1] that took place at the European Semantic Web Conferenece 2006 (ESWC06) this year, as well as the ESWC06 best poster award (*Semantic Wikipedia* – [2]) and the ESWC06 best demonstration award (*semiBlog* – [3]).

This article describes what I believe makes Social Software attractive for the Semantic Web community, and what makes the Semantic Web attractive for the Social Software community. It also derives challenges for the Semantic Web community to address that seem relevant to us based on our experience with Social Software and the Semantic Web. In the remainder of this introduction, I briefly introduce Social Software, the Semantic Web, and the combination of both, which I call *Semantic Social Software*.

Social Software

According to Wikipedia, Social Software is software that "enables people to rendezvous, connect or collaborate through computer-mediated communication and

to form online communities."[1] Although this definition in principle also includes technologies that have already existed for a long time (like email or Usenet), the term Social Software usually only comprises more recent developments like wikis, weblogs, instant messaging (e.g. AIM, ICQ), social bookmarking (e.g. del.icio.us), media sharing (e.g. Flickr, YouTube), and social networking systems (e.g. MySpace, OpenBC).

Today, huge amounts of content are available in Social Software systems. The free web encycolpedia Wikipedia now hosts over 4 million entries of partly astonishing quality. The social networking site MySpace is one of the most popular web sites overall (ranked number 4 by Alexa, following closely after Google, Yahoo, and MSN). According to the weblog index Technorati[2], there are currently about 40 million blogs with a doubling time of about 6 months and around 1.2 million blog posts every day.

What makes Social Software interesting is not only the huge amount of content but that it considerably changes the way content is created and consumed, maybe even more so than the Web did some 15 years ago: where the traditional process of publishing content was expensive and time consuming, Social Software allows virtually everyone to publish on a mouse click. To speak in market terms: with Social Software, *consumers become prosumers*. Because of these radical changes in content production, I consider Social Software a *disruptive technology*.

Semantic Web

The vision of the Semantic Web is to move from "dumb" content that is suitable for presentation only to "smart" content that may be processed by machines and used in different settings. It is also to move from application-centric systems to data-centric systems, and from a Web focussed on information to a Web focussed on relations between things. According to W3C founder and chair Tim Berners-Lee, the Semantic Web will be *the next big thing*.

The current Semantic Web approach may be briefly described as enriching the existing Web with meta-data and (meta-)data processing so as to provide Web-based systems with advanced (so-called intelligent) capabilities, in particular with context-awareness and decision support. What distinguishes the Semantic Web from previous AI approaches is that it assumes a distributed but strongly connected web of small pieces of formal knowledge rather than big, centralised knowledge bases.

Semantic Social Software

Semantic Social Software is the combination of Social Software with Semantic Web technologies. Its basic ideas are on the one hand to improve usage of Social Software by adding metadata and on the other hand to improve the process of

[1] http://en.wikipedia.org/wiki/Social_software
[2] http://www.technorati.com

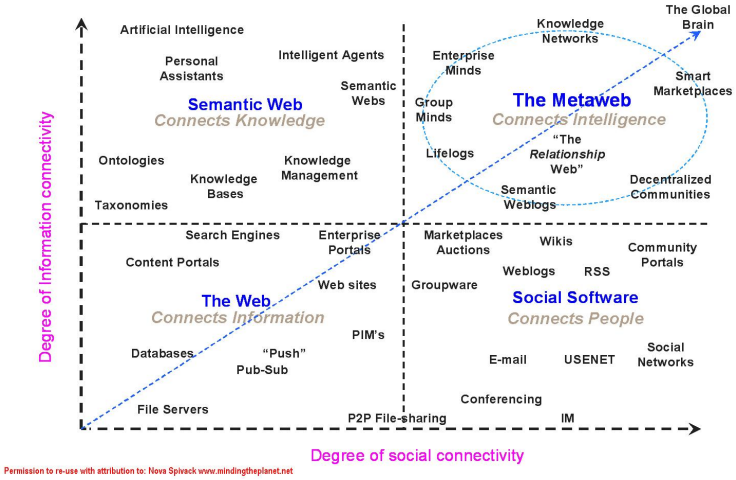

Fig. 1. Nova Spivack's "Metaweb"

creating Semantic Web metadata by using Social Software. These two different perspectives are discussed extensively in the next sections.

The combination of Social Software and Semantic Web technologies seems to fit well. This may be due to the remarkable similarities between Social Software – which is about small but strongly connected pieces of content from different sources with differing opinions – and the Semantic Web – which is about small but strongly connected pieces of formal knowledge from different sources, with different levels of precision and trustworthiness, maybe even inconsistencies. The difference is in this sense only in the level of abstraction: where Social Software mostly deals with social connections and human readable content, the Semantic Web mostly deals with formal connections and formal content. In a way, the step from traditional AI to the Semantic Web is thus very similar to the step from traditional content production to Social Software.

This perspective is also in line with a number of recent predictions for the future of the Web: for example, technology evangelist Nova Spivack outlines in his 2004 articles[3] what he calls "The Metaweb" (cf. Figure 1). The Metaweb is essentially about using social connections to form information connections and vice versa. As can be seen in Figure 1, Nova Spivack is convinced that "The Metaweb is emerging from the convergence of the Web, Social Software and the Semantic Web", connecting human and machine intelligence and moving from "just a bunch of interacting parts" to "a new higher-order whole".

[3] http://novaspivack.typepad.com/nova_spivacks_weblog/2004/03/from_applicatio.html, http://novaspivack.typepad.com/nova_spivacks_weblog/2004/04/new_version_of_.html

This article continues as follows: Section 2 describes two different perspectives one can take on Semantic Social Software, namely the "Semantically Enabled Social Software" perspective (meaning the usage of semantic metadata to enhance existing social software) and the "Socially Enabled Semantic Web" perspective (meaning the usage of Social Software to create semantic metadata). In Section 3, I describe three different kinds of Social Software whose combination with Semantic Web technology appears promising: *Semantic Wikis*, *Semantic Weblogs*, and *ePortfolios*. Out of these three applications, I derive salient aspects of Semantic Social Software in Section 4. I present the mentioned major challenges for the Semantic Web community in Section 5, and conclude with an overview over related work and perspectives.

2 "Semantically Enabled Social Software" or "Socially Enabled Semantic Web"?

Semantic Social Software is traditionally approached from two directions. The first, which I call "Semantically Enabled Social Software", makes use of Semantic Web technology to enhance existing Social Software. The second, which I call "Socially Enabled Semantic Web" makes use of the structures in Social Software to form Semantic Web data. In the following, I describe these two approaches and show that they are really just two sides of the same story, because the difference is primarily in the *use* of applications and not in the applications themselves.

2.1 Semantically Enabled Social Software

As has been lined out in the introduction, massive amounts of digital content are nowadays available in Social Software systems. Although content in Social Software systems is highly connected via hyperlinks and/or social networks, finding relevant content is becoming increasingly difficult, because the existing structure is used merely for presentation purposes. With diverging applications it is furthermore often hard to exchange content between different systems.

"Semantically Enabled Social Software" tries to overcome these issues by applying Semantic Web technology to Social Software. The existing informal or semiformal structures like hyperlinks are augmented by machine-readable formal descriptions ("metadata") that make explicit the actual meaning behind a connection. Such metadata allows for more sophisticated services, like improved search and navigation (e.g. queries on the structure, context visualisation, derived knowledge), personalised presentation of content (e.g. based on personal preferences), and improved interoperability between systems (e.g. to integrate several applications in a company process, or to support smart agents).

For example, a hyperlink from an article in Alice's Weblog to an article in Bobs Weblog could be annotated with "agrees with" or "disagrees with". Articles in a Weblog could furthermore be associated with certain topics like "EU constitution" (as a subtopic of "EU politics"). This would give readers the

opportunity to search not only for articles relevant to a certain topic but also to find supporting or diverging opinions on this topic.

This approach is, among others, followed by the *Semantic Wikipedia* project [2], which aims to enhance the existing Wikipedia by semantic annotations to facilitate search and navigation, and by the *semiBlog* project [4], which aims to simplify search, connectivity, browsing, and management of a Weblog by using semantic annotations.

2.2 Socially Enabled Semantic Web

The "Socially Enabled Semantic Web" sees Semantic Social Software as a tool that simplifies and thus improves the creation of metadata on the Semantic Web, in the same manner as Social Software simplified and improved the creation of data on the normal Web. This approach is founded in the observation that creating formal, machine-readable content on the Semantic Web is hard, which is arguably still one of the most significant barriers to the wide-spread adoption of the Semantic Web. The reason for this is primarily that creating formal metadata currently requires significant expertise in the modeled domain (e.g. biology), as well as the use and intricacies of the used formal languages (e.g. RDF and OWL). The latter issue could be significantly mitigated by appropriate tools, but existing tools are insufficient in the sense that they are much too complicated to use for non-technical users. As a consequence, formalised content on the Semantic Web is still rare and limited to some selected domains.

Semantic Social Software has the potential to overcome some of these problems. It supports the creation of metadata in a number of ways. First, it builds upon existing structures, where connections reflect real-world relationships that are "natural" to the people using the software. Second, it supports the collaboration of people with different backgrounds and expertise, allowing e.g. a biologist and a computer engineer to work on the same knowledge model, augmenting each other. And third, Semantic Social Software provides instant gratification: every bit of formal knowledge contributed by a user is immediately usable.

For instance, a common ontology engineering process supported by a Semantic Wiki (as a special kind of Semantic Social Software) could be to start with a collection of normal Wiki pages (as e.g. found on Wikipedia) that make up the domain to be modeled and then augment the existing hyperlink structure between Wiki pages with machine readable annotations. The first task (writing informal Wiki content) could be easily achieved by a biologist, the second task (formalising the hyperlink structure) by a knowledge engineer. Both could contribute their expertise and collaborate on the creation of the knowledge model.

This approach is e.g. followed by Peter Mika in his work about complementing ontologies with social networks [5] and by our own work on collaborative knowledge engineering with the Semantic Wiki *IkeWiki* [6,7].

2.3 Two Sides of the Same Story

Although the two perspectives described above have originally developed separately and with different application scenarios in mind, the actual software used in both perspectives shares many properties, even to the extent of being actually the same tool used in different settings. This is best exemplified by the history of our own Semantic Wiki system *IkeWiki* [6,7,8]: originally developed to support non-technical domain experts in the process of formalising their knowledge (i.e. "Socially Enabled Semantic Web"), it is now primarily used in a number of settings focussed on enhancing the functionality of Social Software (i.e. "Semantically Enabled Social Software"), like our own group's knowledge base or the conference wiki of the *Social Skills durch Social Software* conference[4]. Similar developments can be seen in other projects.

I therefore argue that, although the two research directions have different goals, they are actually only two sides of the same story, namely "Semantic Social Software". This joint perspective allows to investigate Semantic Social Software as a kind of software that is – based on its properties – helpful in many different settings, rather than looking at it merely as a tool designed for a single purpose. In this article, I present three different applications of Semantic Social Software and discuss them under both, the "Semantically Enabled Social Software" and the "Socially Enabled Semantic Web" perspective.

3 Three Applications of Semantic Social Software

Social Software is very diverse and there exist numerous different applications. In this article, I introduce three applications where a combination with Semantic Web technology is or might be fruitful: *Semantic Wikis* and *Semantic Weblogs* as existing applications, and *ePortfolios* as an emerging application. The convergence of the two different approaches to Semantic Social Software can be seen in all three applications.

3.1 Semantic Wikis: Community-Authored Knowledge Models

A *wiki* is essentially a collection of Web sites connected via hyperlinks. Many different wiki systems exist, but they commonly have a straightforward interface for editing content with a simplified syntax that makes it very easy to set hyperlinks to other pages within the wiki. Therefore, content in a wiki is usually strongly connected via hyperlinks. Furthermore, editing of content in wiki systems is Web-based and access is often unrestricted or at least hardly restricted. Most wiki systems also provide a rollback mechanism for reverting back to previous versions in case of accidental or undesired changes. Wikis are used in many areas, like encyclopedia systems (e.g. Wikipedia), as personal or group knowledge management tools, as collaboration tools, or in collaborative learning environments.

[4] http://eportfolio.salzburgresearch.at/wiki/

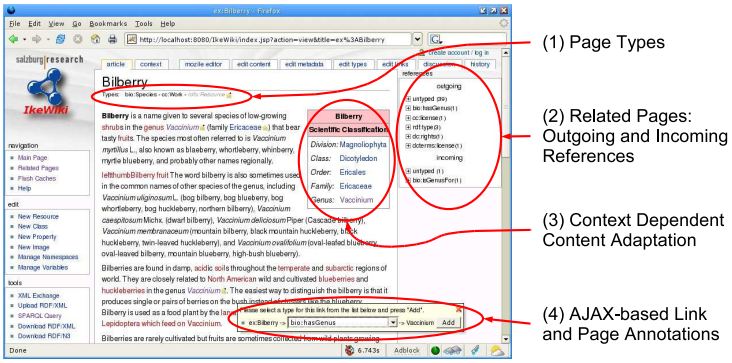

Fig. 2. Page view in IkeWiki interface [7,8]

The main idea of a *Semantic Wiki* is to annotate the inherent hyperlink structure of a Wiki with symbols that describe the meaning of a link in a machine readable fashion. A link from *Mozart* to *Salzburg* could e.g. be annotated with *lived in* or *born in*. Such an annotation can then be used for e.g. enhanced presentation by displaying contextual information, enhanced navigation by giving easy access to relevant related information, and enhanced "semantic" search that allows to query the context in addition to the content.

Semantic Wikis are also excellent tools for collaborative creation of knowledge models. Based on (existing or emerging) natural language descriptions of concepts and individuals, formal knowledge models can be created successively. Natural language descriptions can be created primarily by domain experts (e.g. biologists) and then formalised in collaboration with knowledge engineers. For advanced knowledge engineering tasks, Semantic Wikis offer the possibility to export the knowledge model as RDF or OWL, which can then be loaded into more sophisticated and complicated tools like the OWL ontology editor Protegé.[5]

Semantic Wikis exist in many different flavours (e.g. Semantic MediaWiki [2,9], SemWiki [10], IkeWiki [6,7,8], SemperWiki [11], PlatypusWiki [12]). Some systems are still primarily focused on the page content and see annotations as optional "added value"; they follow more the "Semantically Enabled Social Software" approach (e.g. Semantic MediaWiki). For others, the semantic annotations are in the foreground and sometimes even more important than the actual content; they follow the "Socially Enabled Semantic Web" approach (e.g. IkeWiki and PlatypusWiki).

Figure 2 shows how the page about the *Bilberry* is presented in *IkeWiki*, the Semantic Wiki developed at Salzburg Research. It includes several kinds of semantic metadata: Type information is shown below the page title (1). Links to (semantically) related pages are displayed in a separate "references box" on the right hand side (2). The taxonomy box (3) showing the biological classifi-

[5] http://protege.stanford.edu/

cation of the described plant is automatically generated from existing semantic annotations (i.e. Bilberry *hasGenus* Vaccinium) and is an example for context adaptation and improved functionality by semantic metadata. Finally, (4) shows interactive typing of links using AJAX technology, making it simple for users to add metadata to existing structures.

3.2 Semantic Weblogs: Formalising Discourse

A *weblog* or *blog* is a personal, web-based journal very similar to a digital diary where everyone can read. It is usually organised as a list of short articles ordered by publication date. A very significant difference to a wiki is that content can usually only be authored by the owner of the weblog. Visitors, however, are usually allowed to leave comments about articles that are shown separately. Also, the software used by weblogs is often aware of other weblogs linking to one of the articles and show this information in a so-called *trackback* section as part of the article. Furthermore, weblogs can integrate listings of other weblogs in their navigation (the so-called *blogroll*). In this way, Weblogs often become part of a strongly networked community, the so-called *blogosphere*. Karger and Quan [13] summarise this "essence of blogging" under three key concepts:

1. publishing information in small, discrete notes, as opposed to large, carefully organised Web sites
2. decentralised, per-user publication
3. exposing machine-readable listings

A *Semantic Weblog* uses Semantic Web technology for improved searching and navigation in the blogosphere. Firstly, weblog posts can be associated with formal metadata describing their content, thus improving exchange, search, and retrieval of content. Secondly, the cross-linking between weblog posts in different weblogs can be used for representing the discourse about a certain subject. In [4], the first kind of metadata use is called *content metadata*, while the second kind is called *structure metadata*. Figure 3 (taken from [4]) shows some of the possible structural relations between weblogs.

A particularly interesting property of Semantic Weblogs is the formal representation of discourse about different opinions on a certain topic and the corresponding social networks. This gives readers the possibility to search and navigate not only by topic but also to follow a discussion across different weblogs and to search specifically for certain opinions and connections between weblogs. A salient aspect of discourse representation is the representation of *second/higher order metadata* (or "metadata about metadata"): where normal blog posts may state opinions about the content of other posts, Semantic Weblogs may make formal statements about the formal content of other Semantic Weblogs, requiring more expressive representation formalisms. This could, for example, also be used to form "networks of trust" concerning formal metadata, touching one of the core remaining open issues of the Semantic Web.

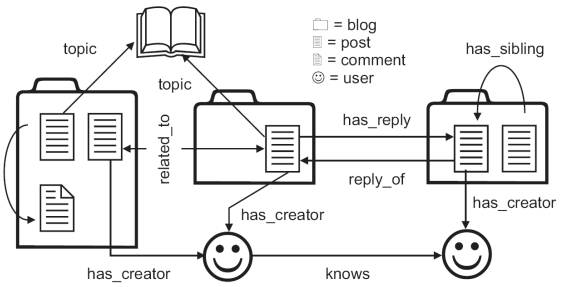

Fig. 3. Structural Relations in the Blogosphere [4]

3.3 ePortfolios: Collecting Learning Artefacts

The *ePortfolio method* [14] is an educational method for supporting life-long learning. The basic idea is to collect all kinds of (digital) artefacts that document a personal learning process in a digital portfolio. Such artefacts may be as coarse as certificates or as fine-grained as every-day documentation of one's learning progress like in a diary, including all relevant material, e.g. test results, articles, presentations, group works, house work, contributions to an online discussion, self-written Wikipedia articles, etc. In other contexts, ePortfolios are for this reason also called *Lifelog* or *Lifeblog*.

The goal of this collection is to be able to reflect about one's own learning progress, to be able to extract parts of one's personal portfolio for external presentation, to design a personal development plan (PDP), to validate and assess the learning outcomes against this plan, and to share learning artefacts with others. Artefacts in such a digital portfolio are strongly connected, both within the personal portfolio and with artefacts in the portfolios of other learners.

Obviously, the ePortfolio method requires software support. An early ePortfolio system is *Elgg* [15].[6] In Elgg, users can sign up and create their own digital portfolio where they can describe their personal skills and interests, publish articles in their personal Weblog, collect links to other websites, and upload files that they consider relevant, particularly newsfeeds from other's portfolios and weblogs. This profile can then, similar to social networking services, be used to find other people with similar interests, or to find communities with interesting topics. In this sense, ePortfolio systems can be seen as *next generation social software*. Other existing ePortfolio systems, like *OSP*[7] or Blackboard's ePortfolio tool, put more emphasis on the integration into learning and course management systems, but the social networking aspect is always existant to a certain degree.

To the best of my knowledge, a *Semantic ePortfolio* does not yet exist (beyond the social tagging used by Elgg). However, as we argue in [16], such software could significantly benefit from Semantic Web technology. Firstly, the collection

[6] http://www.elgg.org
[7] http://www.osportfolio.org/

of artefacts from different systems (like Weblogs, Wikis, online journals) requires interoperability to a degree that is not yet existant. Where current ePortfolio systems can only include raw file content or content created in the systems themselves, a Semantic ePortfolio could go much beyond that and integrate content from different sources and different systems, encoded in different media formats. Secondly, semantic metadata about the artefacts in a digital portfolio could be utilised to more easily create custom-tailored presentations out of the personal portfolio, e.g. for documenting one's learning process for the purpose of evaluation or for creating a tailored resume for the application to a specific job. Thirdly, semantically annotated artefacts could more easily – semi-automatically – be matched against one's personal development plan, allowing for a better reflection by the learner. And fourthly, semantic annotations allow for easier navigation and retrieval through an inevitably huge collection of "life artefacts".

ePortfolios also provide an interesting and challenging application for Semantic Web technologies. As the ePortfolio aims to represent a learner's learning process and thus his skills and knowledge, the semantic metadata in an ePortfolio could mirror the learner's knowledge and interactions with the environment. Such data, if used with the necessary care for privacy, can be useful in many application areas, ranging from automatic creation of personal timetables and development plans over finding relevant jobs in an online job database to assembling a team with complementary skill sets for a certain task in a company.

4 Salient Aspects of Semantic Social Software

Semantic Social Software has a number of properties that make it interesting as a research topic besides the two major "sides" described above. In the following, I give an incomplete overview over what I consider salient aspects of Semantic Social Software.

4.1 Testbed for Semantic Web Technology

Semantic Social Software is software that can be developed quickly and easily, and can build upon existing applications and principles. Augmenting the existing social and hyperlink structures with formal annotations is rather intuitive and a natural extension of existing applications.

At the same time, Semantic Social Software shows many of the promises and also of the problems of the big "Semantic Web Vision". Examples are the improved searching and navigation, personalisation and content adaptation, interoperability, open world assumption, coupling of data and metadata, evolving knowledge models, inconsistencies in real-world data (as many authors work on the knowledge model), ontology alignment (as content from different sources is integrated), etc. Some of these aspects are further elaborated below.

I therefore argue that Semantic Social Software can be an ideal testbed for Semantic Web technology. If problems with the technology arise in Semantic Social Software, they will likely also arise on the large Semantic Web. If technology

works properly in Semantic Social Software, there is also a high chance that it works on the large Semantic Web.

4.2 Coupling of Data and Metadata

An aspect that is often overseen in Semantic Web research is that *metadata* rarely stands for itself. Instead, metadata is in most cases only relevant with respect to the *data* it describes. Combined consideration of data and metadata therefore leads to more realistic settings and "connects human and machine intelligence".

In Semantic Social Software, the combination of data and metadata is quite natural, as existing Social Software with existing content is merely augmented (and not replaced) by metadata. Combined access to metadata and data in Semantic Social Software could e.g. mean enhanced search and navigation or context-dependent presentation of human-readable data.

A combination of data and metadata requires to consider also *query and reasoning languages* that are capable of processing both, data and metadata. An example for such a language is *Xcerpt* [17], a rule-based query language suitable for any kind of semi-structured data, including RDF and HTML. A further issue that often becomes apparent in Semantic Social Software is the so-called *URI crisis*, i.e. the question whether a URI used in the metadata refers to a Web page (Wiki page, Weblog entry) or to the content described in that page (e.g. "elephant").

Another interesting aspect of coupling data and metadata is the semi-automatic extraction of metadata from the actual data using e.g. natural language processing techniques. For example, a page in a Semantic Wiki could be annotated automatically based on the content of the page. As automatic metadata extraction is currently not perfectly reliable, it would make sense to mark such annotations as "automatic" and give users the opportunity to revise them if needed.

4.3 Emerging and Evolving Knowledge Models

A very interesting property of Semantic Social Software is that knowledge models do not appear in discrete "releases", but slowly and dynamically evolve over time. A knowledge model in e.g. a Semantic Wiki may begin with a small set of annotations for a single page to a full-fledged ontology. Also, some parts of the knowledge model will usually be more formalised than others, reflecting areas with more imminent needs or higher community interest.

Such *evolving* knowledge models raise a number of interesting questions that are also relevant for the Semantic Web as a whole (which inevitably will also be an evolving knowledge model). Primarily, evolving systems will be full of *inaccuracies* and even *inconsistencies*, demanding for more tolerant formal languages than those that are proposed today. Also, *trust*, *versioning*, and *merging/aligning* are issues that will need to be addressed appropriately.

4.4 Self-organising Communities Around Emerging Topics

In many Social Software systems, communities organise themselves around certain emerging topics that are relevant at some point in time. In Semantic Social Software, such emerging topics are likely to also be more precisely formalised than others, and could be identified automatically by applying appropriate reasoning. Such information could be used to provide readers with information about "what is relevant", similar to "recent changes" listings in current systems.

5 Challenges for the Semantic Web Community

In our works with and considerations about different kinds of Semantic Social Software, there have been a number of frequently arising challenges that should be taken into consideration by the Semantic Web community if the Semantic Web is supposed to be adopted by a wide community. I present some of these challenges in the following.

"Keep it Simple, Stupid!"

Ultimately, Semantic Web Technology will only "take off" if it is simple enough to use, i.e. the additional effort is in average significantly less than the actual benefit. Systems that require users to learn complicated formal languages imposing (for "outsiders") arbitrary-looking restrictions will very likely fail. Often, as Jim Hendler said,[8] "a little semantics goes a long way." In addition, we need tools that hide most of the complexity of the language from the user, e.g. domain-specific interfaces for creating knowledge and visual interfaces for composing queries. Semantic Social Software goes into this direction, but fails when the underlying languages become too expressive (e.g. number restrictions).

Don't Forget the Data in Metadata

As has been lined out above, metadata is "data about data", and only rarely useful on its own. Considering data and metadata separetely therefore only makes little sense. Instead, it is necessary to develop formalisms, query languages, and reasoning languages that can handle both data and metadata. Steps into this direction have been made (e.g. in Xcerpt [17]), but they are not in the "mainstream" and there is still a long way to go.

Be Tolerant About Inconsistencies

The world is full of inaccuracies and inconsistencies, particularly if different people with different viewpoints work on the same knowledge model – as is the case in Semantic Social Software, and as will be the case in the Semantic

[8] as conference chair in the opening speech of the 2003 International Semantic Web Conference; Sanibel Island, Florida, USA, October 2003

Web. Yet, the Semantic Web, like many or most of its predecessors in AI, builds upon Classsical Logic, which is by the "ex falsum quodlibet" rule intolerant to inconsistencies. What the Semantic Web as a Web of inconsistencies, errors, and inaccuracies needs are formalisms that don't break and can work around problematic parts of the knowledge model (cf. e.g. [18]). Various such formalisms have been proposed in the pre-Semantic Web era, e.g. *paraconsistent logics* or *fuzzy logics*.

There is no "One Size Fits All"

Different applications require different representation formalisms. Whereas languages like OWL are good for representing conceptual hierarchies, they are incredibly bad in representing presentational structures (HTML and CSS do it just fine). What the Semantic Web needs is a flexible and extensible set of formal languages, including formalisms for temporal knowledge, location, rules, situations and events, etc. Even higher-order languages will be useful in many applications, despite their known undecidability.

Reasoning Needs to be Efficient

Many applications work in close interaction with human users. In such situations, it is unacceptable to have complex reasoning systems that take several minutes even for small data sets. It is necessary to improve the performance of existing systems, and to allow developers more fine-grained control over the expressivity and correctness of the reasoner. Often, it is more important to get an almost correct answer within a few milliseconds than to wait for a completely correct answer after several minutes.

Reasoning Needs Truth Maintanance

A significant issue with reasoning systems is that their reasoning is not tracable for users. As a consequence, users often get the feeling of being patronised by the system when the system takes a decision based on reasoning. For Semantic Web applications, it would thus be desirable to allow the user to see all justifications for a decision. A typical example is when the system infers the type for some resource (e.g. a Wiki page) based on the type of a relation to another resource and the user does not agree with this inference, but cannot remove the type directly. In such situations, users should have the option to interfere and modify some of the justifications as needed.

Metadata Needs Versioning

In evolving knowledge models, it is important and interesting to keep track of the modifications done not only in the data but also in the metadata. For this purpose, it is necessary to define what makes up a "unit" or "transaction", as a single modification in the application most likely has effects at many different places in a knowledge model.

6 Related Work

Many people within the Semantic Web community are investigating Semantic Social Software. Most noteworthy is the Semantic Wiki movement that has gained a lot of attention, culminating in this year's 1st Semantic Wiki Workshop [1], colocated with the 3rd European Semantic Web Conference (ESWC06). Other researchers have also worked in the field. An interesting work is the article "Ontologies Are Us" by Peter Mika [5], which discusses the relation between social networks and Semantic Web relationships. Other related works have been mentioned throughout this article, most notably various Semantic Social Software systems, and Nova Spivack's "Metaweb".

Related to the Semantic Web challenges I present in Section 5 is Frank van Harmelen's invited talk entitled *Where does it break? or: Why the Semantic Web is not just "research as usual"* at ESWC06, where he presents similar challenges for the Semantic Web community. Interestingly, I presented our system IkeWiki to him the day before his invited talk, where I – among other things – also discussed with him the challenges presented in this article based on problems we encountered while working with IkeWiki, without knowing that he had very similar issues to present to the community.

7 Conclusion and Perspectives

Semantic Social Software as the combination of Social Software and Semantic Web technologies has great potential, for Social Software developers as well as for Semantic Web researchers. In this article, I presented the two different perspectives on Semantic Social Software and gave three example applications. From our experience with developing and working with Semantic Social Software, I derived challenging research issues for the Semantic Web. These research issues need – in my opinion – to be investigated for making Semantic Web technology successful.

In future work, our group will develop a common framework for implementing Semantic Social Software (and other kinds of semantic software), called the "knowledge-based content management system". The aim of this system is to provide a robust foundation for storing, accessing, and processing content as well as metadata. The system will possibly, and hopefully, solve and implement some of the issues mentioned in Section 5. The combination of Semantic Web technology and Social Software has also been chosen as the next research focus of *Salzburg NewMediaLab*, Austria's industrial competence centre for research on digital content.

Acknowledgements

This work has been partly funded by *Salzburg NewMediaLab*, a competence centre in the K_{ind} action line funded by the Austrian Federal Ministry of Economics and Labor (BMWA) and the state of Salzburg.

References

1. Völkel, M., Schaffert, S., eds.: 1st Workshop "From Wiki to Semantics" (SemWiki'06) – colocated with ESWC'06, Budva, Montenegro (2006)
2. Völkel, M., Krötzsch, M., Vrandecic, D., Haller, H., Studer, R.: Semantic Wikipedia. In: Proceedings of the 3rd European Semantic Web Conference (ESWC06) – Poster Track, Budva, Montenegro (2006)
3. Möller, K.: Publishing Desktop Data with semiBlog. In: Proceedings of the 3rd European Semantic Web Conference (ESWC06) – Demonstrations Track, Budva, Montenegro (2006)
4. Möller, K., Bojars, U., Breslin, J.G.: Using Semantics to Enhance the Blogging Experience. In: Proceedings of the 3rd European Semantic Web Conference (ESWC06). LNCS 4011, Budva, Montenegro, Springer-Verlag (2006)
5. Mika, P.: Ontologies are us: A unified model of social networks and semantics. In: Proceedings of the 4th International Semantic Web Conference (ISWC 2005). LNCS 3729, Galway, Ireland, Springer-Verlag (2005)
6. Schaffert, S., Gruber, A., Westenthaler, R.: A Semantic Wiki for Collaborative Knowledge Formation. In: Semantics 2005, Vienna, Austria, Trauner Verlag (2005)
7. Schaffert, S.: IkeWiki: A Semantic Wiki for Collaborative Knowledge Management. In: 1st International Workshop on Semantic Technologies in Collaborative Applications (STICA'06), Manchester, UK (2006)
8. Schaffert, S., Westenthaler, R., Gruber, A.: IkeWiki: A User-Friendly Semantic Wiki. In: Proceedings of the 3rd European Semantic Web Conference (ESWC06) – Demonstrations Track, Budva, Montenegro (2006)
9. Krötsch, M., Vrandečić, D., Völkel, M.: Wikipedia and the Semantic Web - The Missing Links. In: Proceedings of the WikiMania2005. (2005)
10. Völkel, M., Oren, E.: Personal Knowledge Management with Semantic Wikis. (2006)
11. Oren, E.: SemperWiki: a semantic personal Wiki. In: 1st Workshop on The Semantic Desktop, colocated with ISWC05, Galway, Ireland (2005)
12. Tazzoli, R., Castagna, P., Campanini, S.E.: Towards a Semantic WikiWikiWeb. In: 3rd International Semantic Web Conference (ISWC2004), Hiroshima, Japan (2004)
13. Karger, D.R., Quan, D.: What Would It Mean to Blog on the Semantic Web. In: Proceedings of the 3rd International Semantic Web Conference (ISWC04), Hiroshima, Japan, Springer-Verlag (2004)
14. Barrett, H.C.: Researching Electronic Portfolios and Learner Engagement. Technical report, The REFLECT Initiative (2005)
15. Tosh, D., Werdmuller, B.: Creation of a learning landscape: weblogging and social networking in the context of e-portfolios. Technical report, University of Edinburgh (2004)
16. Hilzensauer, W., Hornung-Prähauser, V., Schaffert, S.: Requirements for Personal Development Planning in ePortfolios supported by Semantic Web Technology. In: 6th International Conference on Knowledge Management (I-KNOW06), Graz, Austria (2006)
17. Schaffert, S.: Xcerpt: A Rule-Based Query and Transformation Language for the Web. PhD thesis, University of Munich (2004)
18. Schaffert, S., Bry, F., Besnard, P., Decker, H., Decker, S., Enguix, C., Herzig, A.: Paraconsistent Reasoning for the Semantic Web (Position Paper). In: Workshop on Uncertainty Reasoning for the Semantic Web (URSW05) at ISWC05, Galway, Ireland (2005)