



Centrum voor Wiskunde en Informatica

**REPORTRAPPORT**

Hypermedia and the Semantic Web: A Research Agenda

J.R. van Ossenbruggen, H.L. Hardman, L. Rutledge

Information Systems (INS)

**INS-R0105 May 31, 2001**

Report INS-R0105  
ISSN 1386-3681

CWI  
P.O. Box 94079  
1090 GB Amsterdam  
The Netherlands

CWI is the National Research Institute for Mathematics and Computer Science. CWI is part of the Stichting Mathematisch Centrum (SMC), the Dutch foundation for promotion of mathematics and computer science and their applications.

SMC is sponsored by the Netherlands Organization for Scientific Research (NWO). CWI is a member of ERCIM, the European Research Consortium for Informatics and Mathematics.

Copyright © Stichting Mathematisch Centrum  
P.O. Box 94079, 1090 GB Amsterdam (NL)  
Kruislaan 413, 1098 SJ Amsterdam (NL)  
Telephone +31 20 592 9333  
Telefax +31 20 592 4199

# Hypermedia and the Semantic Web: A Research Agenda

Jacco van Ossenbruggen, Lynda Hardman and Lloyd Rutledge

*Firstname.{van.}Lastname@cwi.nl*

CWI

*P.O. Box 94079, 1090 GB Amsterdam, The Netherlands*

## ABSTRACT

Until recently, the Semantic Web was little more than a name for the next generation Web infrastructure as envisioned by its inventor, Tim Berners-Lee. Now, with the introduction of XML and RDF, and new developments such as RDF Schema and DAML+OIL, the Semantic Web is rapidly taking shape. In this paper, we first give an overview of the state-of-the-art in Semantic Web technology, the key relationships with traditional hypermedia research, and a comprehensive reference list to various sets of literature (Hypertext, Web and Semantic Web). The core of the paper presents a research agenda by describing the open research issues in the development of the Semantic Web from the perspective of hypermedia research.

*1998 ACM Computing Classification System:* H.5.4, H.5.1, I.2.4

*Keywords and Phrases:* Hypermedia, Semantic Web, XML, RDF, DAML+OIL, SMIL.

*Note:* The research reported here has been carried out under the "ToKeN2000" project.

## 1. INTRODUCTION

The bulk of content that is currently available on the Web is notoriously hard to process automatically: "... data transmitted across the Web is largely throw-away data that looks good but has little structure" [20]. Markup languages such as (X)HTML [68], SVG [36] and SMIL [65] are primarily geared to documents whose content should be interpretable by *human* interpreters, and hence tend to focus primarily on document structure and document presentation. Little or no attention is given to the representation of the semantics of the content itself, i.e. the (domain-specific) representation of the subject of the document.

In contrast, knowledge representation techniques developed within the Artificial Intelligence (AI) community have a strong tradition in describing domain-specific knowledge in a machine-processable manner. In addition, the digital library community has studied issues related to more persistent ways of storing and cataloging digital content. Recently, initiatives within and outside the World Wide Web Consortium (W3C) are building upon the expertise of these communities by developing knowledge representation and annotation languages on top of the current Web infrastructure. This not only allows newly encoded knowledge to be easily disseminated over the Web, but also provides a convenient syntax for annotating existing content, such as (X)HTML or SMIL content. This combination is a key enabler to the main objective of the Semantic Web: documents whose content is processable by both humans and machines.

While the Semantic Web appears at first sight to be far from the current research trends of the hypertext community, much earlier work in the field lay extremely close to the borders of knowledge representation, for example [18, 19, 47, 51, 59]. These authors were attempting to bridge the gap between knowledge representation and information presentation in a technological context that lacked support for this integration. The Web today provides a sound technological basis for document processing and already supports the first layers of the Semantic Web. In this paper we briefly sketch the current developments of the Semantic Web, compare these with the issues long ago fielded in the hypertext literature and highlight those that should form the basis of a research agenda for a universal




Figure 1: Document and knowledge representation languages on the Web.

information repository.

## 2. THE CURRENT SEMANTIC WEB INFRASTRUCTURE

Figure 1 provides an overview of both the document and knowledge representation languages on the Web. Following current document languages such as XHTML, SVG and SMIL (in the left half of the figure), the various layers of the Semantic Web are all built on top of XML [9], as shown in the right half of the figure. This makes generic XML-based software and languages such as XML parsers, transformation engines (XSLT [16]), path and pointer engines (XPath, XPointer [17, 30]), style engines and formatters (CSS, XSL [8, 26]), etc. directly available on the Semantic Web.

### 2.1 RDF and RDF Schema

The second layer of the Semantic Web infrastructure is the Resource Description Framework (RDF [66]). RDF provides a simple data model for expressing statements using (*subject, predicate, object*) triples, and an associated serialization syntax in XML. All three elements of the triple can be defined within the current document or refer to another resource on the Web. To make statements about a collection of resources, RDF specifies a simple container model, modeling sequences (ordered), bags (unordered) and lists of alternatives. RDF also supports *reification*, that is, statements about other RDF statements.

A set of RDF statements uses a particular vocabulary that defines the properties and data types that are meaningful for the application at hand. Such an RDF vocabulary can be defined by using RDF Schema (RDF-S [67]). As part of its schema language, RDF-S also defines some predefined concepts, including primitives to model a class/subclass hierarchy, relationships between classes ("properties"), and domain/range restrictions on such properties. Note that while the RDF model by itself merely provides a set of triples, RDF-S is already sufficiently expressive to describe a class hierarchy which allows some useful querying and reasoning support. For example, one could query an RDF-S system whether a given instance belongs to a specific class, what (inherited) properties it has, etc [45].

### 2.2 DAML+OIL

While several applications are built directly on RDF and RDF-S, another interesting layer (currently under development) is DAML+OIL [28, 62]. RDF-S is missing some features that are commonly found in systems developed within the AI community (e.g. frame-based systems), while it also contains some features (most notably reification) that make it hard to provide a formal semantics for RDF-S and to

provide fully automated and efficient inference engines.

DAML+OIL addresses these issues by removing support for reification, and extending RDF-S with concepts commonly found in frame-based languages. The result is a language that is compliant with RDF and RDF-S, has a sound formal semantics and an efficiently implemented inference engine. This allows not only more advanced quering, but the inference engine can also be used to detect contradictions and other errors in a DAML+OIL specification.

### 2.3 Applications: PICS, P3P, Dublin Core

Examples of applications that use the infrastructure sketched above include W3C's Platform for Internet Content Selection (PICS [52]), Platform for Privacy Preferences Project (P3P [22]) and the Dublin Core [15]. While PICS was defined before its more generic successor RDF, a mapping to RDF has been developed [10]. Dublin Core also predates RDF, but now also has an RDF-based serialization syntax.

## 3. RELATION WITH HYPERMEDIA RESEARCH

While the Semantic Web aims primarily at providing a generic infrastructure for machine-processable Web content, it has direct relevance to hypermedia research. In order to capture the breadth of relevance of the Semantic Web to hypermedia research, we have analyzed the visionary articles of Malcolm et al. [49], Engelbart [35] and Halasz [37]. A large proportion of these features relate directly to the Semantic Web. On the one hand, the Semantic Web infrastructure should enable several features commonly found in systems developed within the hypermedia community that are currently missing on the Web. On the other hand, the development of the currently emerging Semantic Web infrastructure could directly benefit from the models, systems and lessons learned within the hypermedia community.

Based on the articles mentioned above, we identified around thirty features which have been grouped into the seven categories discussed below:

1. **Basic nodes, links and anchor data model** — Many hypermedia systems feature a model that is similar to the typical data model of nodes, links and anchors defined by the Dexter Hypertext Reference Model [38]. This model is directly applicable to the Semantic Web. To be able to annotate a specific portion of a Web resource, it needs an anchoring mechanism, and to establish a relationship between the annotation and the target resource, a linking model is necessary. The remaining features discussed below can be seen as variations on, or applications of, this basic model.
2. **Conceptual hypertext** — Conceptual hypertext systems introduced a layered hypermedia model, adding a hyperlinked network of related index terms (or concepts) on top of a hyperlinked document base. Additional links up and down between the two levels relate the information in the documents to the concepts in the hyperindex [13, 18]. More recent approaches, such as COHSE [14], go even further and use the full power of *ontologies* to improve hypertext linking based on the semantic relations among the associated concepts.
3. **Typed nodes, links and anchors** — Many hypertext systems base a large part of their functionality on their ability to assign *types* to nodes [44], links [60], and to a lesser extent, anchors [54]. Argumentation systems such as gIBIS [19], for example, use link types to label “response-to” or “object-to” relationships (note that such relationships may, but need not be, represented by a navigational hyperlink in the user interface). RDF allows embedded and external annotation of links and anchors, and with languages such as RDF-S and DAML+OIL, one can easily define an (extensible) type system for links and anchors. For example, RDF-S allows “object-to” to be defined as a subtype of “response-to”, and in DAML+OIL one could define “is-criticized-by” as the inverse of an “object-to” relation.
4. **Virtual links and anchors** — Systems such as Microcosm [24, 41] featured virtual (or “dynamic”) links and anchors. That is, they support run-time computation of links and anchors in

addition to statically defined links and anchors that are defined at authoring time. While the current Semantic Web developments tend to be mainly language-oriented (standard interfaces for generic RDF(S)-based services are yet to be defined), an RDF(S) query/inferences engine could provide an excellent basis for semantically driven hyperlink services. Related areas include ontology-driven linking as discussed in [21] and agent-based navigation assistance as discussed in [33].

5. **Searching and querying** — The need for supporting good search and query interfaces was recognized by the hypermedia community long before the appearance of the first search engines on the Web. In one of his famous “Seven Issues”, Halasz already explained the need for both content-based and structure-based retrieval on hypertexts [37]. In addition, the digital library community has always stressed the use of cataloging techniques and metadata-based search [31]. While this has still to be proven in practice, RDF-enabled search engines have the potential to provide a significant improvement over the current keyword-based engines, especially when it comes to metadata and structure-based searching. An example of such a system, albeit not using RDF for encoding its semantic annotation is the Ontobroker system discussed in [27]
6. **Versioning and authentication features** — While features such as versioning, concurrency, and authentication are not commonly recognized as fundamental hypermedia features, they have frequently been topics of hypermedia research because they are essential for one of the most important hypermedia application domains: Computer Supported Collaborative Work (CSCW). CSCW has been, for example, the driving force for most of Engelbart’s work on NLS/Augment [34, 35] and is listed as one of Halasz’s seven issues [37], etc. Research on CSCW has also been carried out in the context of hypermedia systems such as NoteCards [61] and CHIPS [69].

Because early generations of hypermedia systems were designed as stand-alone systems or as part of an organization’s local network, these features are even more important in Web-based collaboration. It is only because of the Web’s initial focus on “read-only” browsing that these features hardly received any attention. A notable exception is the joint IETF/ W3C work on WebDAV [32]. While WebDAV predates RDF, it has a similar property-based model for Web resources.

7. **Annotation** — The ability to annotate the work of others has traditionally been an important feature of many hypertext systems, and it is another key feature of collaborative hypermedia systems. Annotation has been lost on the Web, however, due to HTML’s embedded link syntax. This syntax requires a user to have write access to the original page to be able to annotate it, which is hardly a realistic requirement on the Web. RDF and its relatives are designed to make statements about any resource on the Web (that is, anything that has a URI), without the need to modify the resource itself. This allows for rich annotations and encoding of semantic relationships among resources on the Web.

Despite the many relations between the Semantic Web and previous hypermedia research, many new research questions arise. To provide the reader with an intuitive feel for the range of issues, the following section describes an example Semantic Web application in the domain of adaptive hypermedia.

#### 4. MOTIVATING EXAMPLE: ADAPTIVE HYPERMEDIA PRESENTATION GENERATION

Our current research focuses on the automatic generation of time-based hypermedia presentations to adapt to the specific preferences and capabilities of individual users and the device they are using. We have implemented a prototype hypermedia generation engine that is able to transform a high-level description of a hypermedia presentation into a concrete hypermedia final-form encoding that is readily playable on the end-user’s system [57]. Note that in the context of this paper, we are




Figure 2: Information sources used by the Cuypers adaptive hypermedia generation engine.

primarily interested in the nature of the information the system uses as its input, and the information it produces. The inner details of the Cuypers system itself have been discussed in previous work [64].

Adaptive hypermedia systems such as Cuypers adapt hypermedia presentations to the specific context of an individual user [7, 11]. This involves semantic modeling on several levels. First, an adaptive hypermedia system needs an adequate user profile that models the preferences and capabilities of the user [12]. Second, it needs a model of the domain at hand, and third, a model of how, given a user profile and a domain model, information about this domain should be conveyed to the user (cf. the “teaching model” discussed by De Bra et al. in [25]). Fourth, a system-oriented profile is needed that models the network resources and presentation capabilities of the user’s platform (e.g. the amount of available bandwidth, what file formats are recognized etc.). Finally, the system needs a model that describes how a given presentation can be best adapted to the target platform [6, 64].

While all these different types of knowledge are already used within the current Cuypers prototype (as depicted in Figure 2), it is either encoded using ad-hoc encoding representation techniques or, more often, remains implicit and hidden in the (procedural) generation software. In addition, the prototype simply assumes that all relevant knowledge about its users and their systems is available, which is in practice rather naive in the context of privacy concerns. In general, the large amount of modeling needed in most adaptive hypermedia systems is a major barrier to a wider adoption of this technology.

The Semantic Web can — at least partially — solve these issues. First, it provides a declarative, interoperable foundation for modeling, encoding, dissemination and processing these different types of knowledge in a common, interoperable way while respecting the privacy of all parties involved. For example, the requirement to be able to describe the wide variety of capabilities that exists among mobile phones has led to the more general CC/PP initiative [55] that aims at providing a standardized (RDF) vocabulary [46] for describing the characteristics of a Web client application. When vendors make CC/PP profiles of their products publicly available (either as part of the product or via their Web sites), adaptive systems, such as Cuypers, could simply reuse this information, and no longer need to develop their own profiles of the platforms they target.

As another example, the system could use the P3P framework [22] to automatically negotiate with

```

<smil xmlns="http://www.w3.org/2000/SMIL20/CR">
  <head>
    <meta name="generator" content="CWI/Cuyppers 1.0"/>
    <metadata>
      <rdf:RDF xml:lang="en"
        xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
        xmlns:oil="http://www.ontoknowledge.org/oil/rdf-schema/2000/11/10-oil-standard"
        xmlns:museum="http://ics.forth.gr/.../museum.rdf"
        xmlns:token="http://www.token2000.nl/ontologies/additions.rdf">
        <rdf:Property rdf:about="http://www.token2000.nl/additions.rdf#painted-by">
          <oil:inverseRelationOf
            rdf:resource="http://ics.forth.gr/.../museum.rdf#paints"/>
        </rdf:Property>

        <museum:Museum rdf:ID="Rijksmuseum"/>
        <museum:Painter rdf:ID="Rembrandt">
          <museum:fname>Rembrandt</museum:fname>
          <museum:lname>Harmenszoon van Rijn</museum:lname>
        </museum:Painter>
        <museum:Painting rdf:about="#apostlePaul">
          <museum:exhibited rdf:resource="#Rijksmuseum" />
          <museum:technique>chiaroscuro</museum:technique>
          <token:painted-by rdf:resource="#Rembrandt" />
        </museum:Painting>
      </rdf:RDF>
    </metadata>
    <layout>...</layout>
  </head>
  ...

```

Figure 3: Example of annotated SMIL using OIL and RDF-S (header).

the user's client system about the user's personal profile. It could, for example, try to get more information from the user by declaring that the information will only be used for adaptation in the current session and will not be permanently stored nor used for other purposes.

Similar arguments for reusing information available on the Semantic Web apply to the other types of knowledge used by the system. In addition to using the Semantic Web on the input side, the system could also use it in the encoding of the hypermedia presentation it generates as its output. We are currently investigating the automatic generation of richly annotated SMIL documents in the context of a museum application. Figure 3 shows (part of) the annotations embedded in the header of an example SMIL presentation (shown in Figure 5 on page 8).

Note that the nature of such annotations may depend on the user's context, that is, the annotations themselves are also subject to adaptation. This particular example is taken from a scenario where the user asks the system about the use of the *chiaroscuro* technique (that is, the use of high contrast lighting) in the paintings of Rembrandt. Here, the annotations define the *Rijksmuseum* as being an instance of the *Museum* class, define *Rembrandt* as being a *Painter* and define the element `apostlePaul` (defined in the body of the document, see Figure 4 on the facing page) as a *Painting* that is exhibited in the *Rijksmuseum*, that uses the technique *chiaroscuro* and is painted by Rembrandt. The current Cuyppers prototype generates a hypermedia presentation in SMIL that gives a textual description of the term *chiaroscuro*, illustrated by a sequence of Rembrandt paintings that are annotated as being examples of this technique. We are currently extending the system so that it can also reproduce the relevant semantics in the form of an annotated SMIL document, as shown in Figure 3.



```

...
<body>
  <par>
    <text region="title" src="...queries to the multimedia database..."/>
    <text region="descr" src="...to find descripton of 'chiaroscuro'..."/>
    <seq>
      <par dur="10" id="apostlePaul">
        
        <text region="ptitle" src="...the title of the painting.. "/>
      </par>
      <par dur="10"> ... 2nd painting+title ... </par>
      <par dur="10"> ... 3rd painting+title ... </par>
      ...
    </seq>
  </par>
</body>
</smil>

```

Figure 4: Example of annotated SMIL using OIL and RDF-S (body).

Note that, if appropriate, the system could also generate annotations on the basis of the other knowledge sources it has access to. For example, it might also annotate the media elements in terms of their role in the discourse model (by stating that a piece of a text is a *definition* of a particular concept and that an image is an *example* of this concept).

Such a highly adaptive notion of semantic annotation requires a flexible annotation syntax which can make use of multiple vocabularies, as shown by the example document. In the example, the SMIL namespace is the document's default namespace, but the included annotations combine concepts from a variety of other namespaces (as shown by the `xmlns` declarations in the beginning). Note that while some common concepts are taken directly from RDF or OIL, most concepts are defined by a domain-specific ontology. For this example, we did not need to develop our own ontology for the museum domain, but simple reused the ontology discussed in [45]. Adding additional annotations such as *definition* and *examples* would simply require the use of yet another RDF vocabulary.

There are, however, also a number of serious limitations to the techniques used in the example.

First of all, while the semantics of the annotations in the example may seem straightforward at first sight, in the general case it remains unclear what the semantics are of RDF vocabularies that are mixed and refer to one another on an ad-hoc basis.

Another limitation is the embedded encoding of the annotations in the example. While RDF allows encoding of the annotations separately from the SMIL document, this would require new annotation services with the functionality similar to that of open hypermedia systems. If, for example, we were to store the annotations separately (e.g. on another site) from the SMIL presentation, there would be no common infrastructure to serve the annotations to the client applications that process the SMIL presentation.

A third problem is related to the anchoring mechanism used. The annotations about the painting could conveniently use a standard URI fragment identifier to point to the target element labeled `apostlePaul`. In a more realistic decription environment, such as that described in [53], descriptions would be associated with different fragments of the different (time-sensitive) media types. When annotations need to point to a fragment of an image, video or other non-XML media item, we run into problems because there is no commonly agreed upon fragment identifier format (i.e. no anchoring mechanism) for non-XML media [63].

The Web is all about sharing information, and the Semantic Web is no different. Even in our simple example we have already used two different ontologies. In addition to the problems related

to the semantics of merging and linking ontologies, this also illustrates problems of sharing existing ontologies, collaborative development of new ontologies and of assuming the level of trustworthiness of ontologies created by others.

The example given in this section aimed at providing an intuitive notion of the possibilities and limitations of the current Semantic Web infrastructure. In the following section, we provide a more systematic discussion of the open research questions and the potential contribution of the hypermedia research community in resolving them.



Figure 5: Example SMIL presentation about the use of *chiaroscuro* in the works of Rembrandt.

## 5. OPEN RESEARCH QUESTIONS

Before the true potential of the Semantic Web can be fully exploited, a number of key issues need to be resolved. In this section we identify open issues related to links and relationships, open hypermedia, time-based hypermedia and computer-supported collaborative work.

### 5.1 *Links versus Relationships*

While the current Semantic Web languages are strong in representing (semantic) relationships between Web resources, this is insufficient for full hyperlink support. First, in addition to the currently defined *languages*, hypermedia applications also need to be able to access the associated *services*. For example, given an RDF annotation, finding the resources this annotation is about is simply a matter of dereferencing the URIs used. The other way round, however, is a lot harder. This requires intranet or even Internet crawlers that collect and index RDF annotations so that, given a particular Web resource, one can find the relevant annotations associated with that resource (the issues related to the software architecture of such services are discussed in the “Open hypermedia and the Semantic Web” section below).

Another issue is the fact that the Web uses different approaches for modeling and encoding links and relationships across Web resources. In addition to the RDF family discussed above and the embedded links commonly found in Web languages such as HTML, WML and SMIL, W3C is also developing the XML Linking Language (XLink [29]) as a common syntax for encoding embedded and non-embedded links in XML documents. When compared to RDF, XLink provides some extra built-in link functionality (some basic traversal behavior, for example). XLink’s ability to encode semantic

relationships, however, is far less than that of RDF, and XLink's hyperlink syntax is not backward-compatible with that of HTML, WML or SMIL. Whether the extra link functionality of XLink is sufficient to justify widespread adoption is still a matter of debate. Sticking to HTML's approach for simple, embedded links while adopting the full power of the RDF family for encoding extended and external links seems to be a viable alternative. For example, taxonomic hypertext systems might benefit more from ontology-oriented languages such as DAML+OIL than from languages oriented towards navigational hyperlinks such as XLink.

A third, and more complex, issue we want to discuss is not related to linking across documents, but to linking across knowledge sources. Traditionally, knowledge bases, expert systems, ontologies, etc., as developed within the AI community, focus on representing centralized, consistent and trustworthy knowledge. On the Web, knowledge is typically decentralized, inconsistent and not always to be trusted. These differences bring up new, fundamental problems, most of which remain to be solved. For example, most of the problems that arise when linking in fragments from one ontology into another are still unsolved. On the Web, an application has to be able to deal with distributed, cross-linked, incompatible or even inconsistent pieces of knowledge. A related issue is the requirement to be able to use terms from different ontology fragments. For example, Hunter et al. [42] describe the issues that arise when multiple metadata ontologies need to be used within a single application profile.

### 5.2 Open hypermedia and the Semantic Web

Open hypermedia systems (OHS) aim at adding hypermedia functionality to existing applications with minimal impact on the original application and its native file format [56]. These goals explain two fundamental differences between the OHS approach and the Web. First, while the majority of the links on the Web are *embedded links*, OHS focus on encoding links externally from the documents being linked, in order to preserve the application's native file format. Second, while Web browsers implement linking functionality within the browser, OHS architectures require minimal extra functionality of the client application because most of the link services are realized by a dedicated link server.

While the reduced complexity of embedded links on the Web has many advantages [70], for the Semantic Web the OHS approach seems more realistic. First of all, the traditional "to embed or not to embed" discussion [23] also applies to the Semantic Web. Semantic relationships are, even on the Web, expected to be significantly more complex than HTML's simple, uni-directional links. Embedded encoding of such information will increase the complexity of authoring Web content and increase maintenance costs when keeping Web pages up-to-date. In addition, bulky annotations will increase downloading times for all applications, even those that do not need to (or cannot) process the semantic annotations. In addition, the processing of (domain-specific) semantic annotations is likely to be domain specific in itself, and will thus vary from site to site. Implementing specific reasoning and inference services makes sense only at the server-side and not in a generic Web-client. The picture sketched above, with a focus on externally encoded semantic relationships and dedicated server applications to maintain and process these semantics is very similar to the OHS approach. It suggests that many of the lessons learned in OHS modeling, software architecture and the design of interoperable protocols will be directly applicable to the Semantic Web. Especially in the context of the current, mainly language-driven, developments on the Web, open hypermedia systems may very well provide a blueprint for an emerging Semantic Web infrastructure. Such an infrastructure should provide interoperable interfaces and protocols to a variety of annotation services. Examples of such services include the common storage, maintenance and retrieval of semantic annotations on the Web, and the (domain) specific reasoning and inference engines that use these data effectively.

### 5.3 Time-based hypermedia and the Semantic Web

Time-based hypermedia systems integrate hyperlink navigation with synchronized multimedia presentation [39]. They bring with them problems of timing and synchronization, inclusion of different media and streaming of data-intensive media, such as video and audio. Time often plays an important role, on multiple levels, in the modeling of the semantics, narrative and document structure of hy-

permedia content [40, 48, 58]. The special role of time, and also space [50], in describing hypermedia content and hypermedia structure seems to justify the representation of these concepts as primitives of standardized hypermedia annotation vocabularies that could be build on top of languages such as RDF-S and DAML+OIL.

To integrate time-based hypermedia into the Semantic Web, a requirement is that we are able to annotate multimedia content as easily as text-based (XML) content. Because existing pointing languages such as XPath and XPointer are limited to XML content, new languages need to be developed to be able to point into the time-variant, binary encoded and compressed data formats that are common in the multimedia domain. To optimize both the quality of the presentation as well as the interactive response times, streamed delivery of media content is currently the norm in distributed multimedia environments such as the Web. Downloading bulky metadata in today's non-streamable formats is a major threat to both presentation quality and interactive response time. Instead, we need to investigate streamable versions of the RDF family of languages, and — probably even harder — the associated (incremental) reasoning and inference algorithms.

#### *5.4 CSCW and the Semantic Web*

Even with the current Semantic Web infrastructure and distributed Web authoring protocols such as WebDAV, many of the features related to authentication, access control, concurrency control and version control as discussed by [35, 37, 49] are not yet fully integrated in the Web's infrastructure. Part of this problem could be addressed by providing interoperable realizations of these features in the form of extensions to and layers upon the currently available protocols. This would, however, only solve the technical part and neglect the social and dynamic aspects of collaboration. Addressing this part of the problem requires integration of the Semantic Web infrastructure into collaborative tools that support typical groupware features related to awareness, synchronous and asynchronous communication and workflow-oriented systems that explicitly support dependencies between user tasks and other coordination mechanisms.

## 6. CONCLUSION

In this paper, we gave an overview of the developing Semantic Web infrastructure, showed how this relates to typical hypermedia research topics and gave comprehensive pointers to the relevant literature. We used the generation of annotated hypermedia presentations as an example scenario to illustrate the potential applications of the Semantic Web and to exemplify the related research questions that still need to be solved. We then described four areas of research that need to be addressed in order to allow the Semantic Web to realise its full potential.

Originally, hypertext research aimed at bringing the user's interaction process with digitally stored information closer to the semantic relations implicit within the information. Much of the more 'hypertext-specific' research, however, turned to system and application-oriented topics, possibly through the lack of an available infrastructure for supporting more explicit semantics. The introduction of the Web, as a highly distributed, but relatively simple, hypermedia system, has also influenced the character of hypermedia research. The existence of XML and RDF, along with developments such as RDF Schema and DAML+OIL, are giving the impetus for realizing the Semantic Web. During these early stages of its development, we want to ensure that the many hypertext lessons learned in the past will not be lost, and that future research tackles the most urgent issues of the Semantic Web.

#### ACKNOWLEDGEMENTS

Frank van Harmelen provided many insightful comments on a preliminary draft of this paper. Part of the research described here has been carried out in the context of the Eureka/ITEA project RTIPA and two national projects funded by the Dutch government: Dynamo and Token2000.

## References

1. ACM. *Hypertext '87 Proceedings* (Chaper Hill, North Carolina, November 13-15, 1987), ACM Press.
2. ACM. *Proceedings of the ACM European Conference on Hypermedia Technology (ECHT'94)* (September 18-23, 1994, Edinburgh, 1994), ACM Press.
3. ACM. *The Proceedings of the Ninth ACM Conference on Hypertext and Hypermedia* (Pittsburgh, PA, June 20-24, 1998), ACM Press. Edited by Kaj Grønbaeck, Elli Mylonas and Frank M. Shipman III.
4. ACM. *Proceedings of the 10th ACM conference on Hypertext and Hypermedia* (Darmstadt, Germany, February 21-25, 1999). Edited by Klaus Tochtermann, Jorg Westbomke, Uffe K. Will and John J. Leggett.
5. ACM. *Proceedings of the 11th ACM conference on Hypertext and Hypermedia* (San Antonio, Texas, USA, May 30 – June 3, 2000).
6. BOLL, S., KLAS, W., AND WANDEL, J. A Cross-Media Adaptation Strategy for Multimedia Presentations. In *ACM Multimedia '99 Proceedings* (Orlando, Florida, October 30 - November 5, 1999), ACM, Addison Wesley Longman, pp. 37-46.
7. BORDEGONI, M., FACONTI, G., MAYBURY, M., RIST, T., RUGGIERI, S., TRAHANIAS, P., AND WILSON, M. A Standard Reference Model for Intelligent Multimedia Presentation Systems. *Computer Standards & Interfaces* 18, 6-7 (December 1997), 477-496.
8. BOS, B., LIE, H. W., LILLEY, C., AND JACOBS, I. Cascading Style Sheets, level 2 CSS2 Specification. W3C Recommendations are available at <http://www.w3.org/TR>, May 12, 1998.
9. BRAY, T., PAOLI, J., AND SPERBERG-MCQUEEN, C. M. Extensible Markup Language (XML) 1.0 Specification, February 10, 1998. W3C Recommendations are available at <http://www.w3.org/TR>.
10. BRICKLEY, D., AND SWICK, R. R. PICS Rating Vocabularies in XML/RDF. W3C Note for discussion only. Available at <http://www.w3.org/TR/rdf-pics>, 27 March 2000.
11. BRUSILOVSKY, P. Methods and Techniques of Adaptive Hypermedia. *Journal on User Modeling and User-Adapted Interaction* 6 (1996), 87-129.
12. BRUSILOVSKY, P., KOBSA, A., AND VASSILEVA, J., Eds. *Adaptive Hypertext and Hypermedia*. Kluwer Academic Publishers, Dordrecht, The Netherlands, 1998.
13. BRUZA, P. D. Hyperindices: A Novel Aid For Searching in Hypermedia. In *Hypertext: Concepts, Systems and Applications, Proceedings of the European Conference on Hypertext* (INRIA, France, November 1990), Cambridge University Press, pp. 109-122.
14. CARR, L., BECHHOFFER, S., GOBLE, C., AND HALL, W. Conceptual Linking: Ontology-based Open Hypermedia. In *The Tenth International World Wide Web Conference* [43], pp. 334-342.
15. CATHRO, W. Metadata: An Overview. In *Standards Australia Seminar: Matching Discovery and*

- Recovery* (August 1997).
16. CLARK, J. XSL Transformations (XSLT) Version 1.0. W3C Recommendations are available at <http://www.w3.org/TR/>, 16 November 1999.
  17. CLARK, J., AND DEROSE, S. XML Path Language (XPath) Version 1.0. W3C Recommendations are available at <http://www.w3.org/TR/>, 16 November 1999.
  18. COLLIER, G. H. Thoth-II: Hypertext with Explicit Semantics. In *Hypertext '87 Proceedings* [1], pp. 269–289.
  19. CONKLIN, J., AND BEGEMAN, M. L. gIBIS: A Hypertext Tool for Team Design Deliberation. In *Hypertext '87 Proceedings* [1], pp. 247–251.
  20. CONNOLLY, D. W. An Evaluation of the World Wide Web with respect to Engelbart's Requirements. <http://www.w3.org/Architecture/NOTE-ioh-arch>, February 1998.
  21. CRAMPES, M., AND RANWEZ, S. Ontology-supported and ontology-driven conceptual navigation on the World Wide Web. In *Proceedings of the 11th ACM conference on Hypertext and Hypermedia* [5], pp. 191–199.
  22. CRANOR, L., LANGHEINRICH, M., MARCHIORI, M., PRESLER-MARSHALL, M., AND REAGLE, J. The Platform for Privacy Preferences 1.0 (P3P1.0) Specification. W3C Candidate Recommendations are available at <http://www.w3.org/TR/>, 15 December 2000. Edited by Massimo Marchiori.
  23. DAVIS, H. To Embed or Not to Embed... *Communications of the ACM* 38, 8 (August 1995), 108–109.
  24. DAVIS, H., HALL, W., HEATH, I., HILL, G., AND WILKINS, R. Towards an Integrated Information Environment with Open Hypermedia Systems. In *Proceedings of the ACM European Conference on Hypertext (ECHT'92)* (Milano, Italy, November 30 - December 4, 1992), ACM, ACM Press, pp. 181–190.
  25. DE BRA, P., HOUBEN, G.-J., AND WU, H. AHAM: A Dexter-based Reference Model for Adaptive Hypermedia. In *Proceedings of the 10th ACM conference on Hypertext and Hypermedia* [4], pp. 147–156. Edited by Klaus Tochtermann, Jorg Westbomke, Uffe K. Will and John J. Leggett.
  26. DEACH, S. Extensible Stylesheet Language (XSL) Specification. W3C Candidate Recommendations are available at <http://www.w3.org/TR/>, 21 November 2000.
  27. DECKER, S., ERDMANN, M., FENSEL, D., AND STUDER, R. Ontobroker: Ontology Based Access to Distributed and Semi-Structured Information. In *Semantic Issues in Multimedia Systems. Proceedings of DS-8 (1999)*, R. M. et al, Ed., Kluwer Academic Publisher, pp. 351–369.
  28. DECKER, S., MELNIK, S., HARMELEN, F. V., FENSEL, D., KLEIN, M., BROEKSTRA, J., ERDMANN, M., AND HORROCKS, I. The Semantic Web: The roles of XML and RDF. *IEEE Internet Computing* 15, 3 (October 2000), 63–74.
  29. DEROSE, S., MALER, E., AND ORCHARD, D. XML Linking Language (XLink). W3C Proposed Recommendations are available at <http://www.w3.org/TR/>, 20 December 2000.
  30. DEROSE, S., MALER, E., AND RON DANIEL, J. XML Pointer Language (XPointer) Version 1.0. W3C Candidate Recommendations are available at <http://www.w3.org/TR/>, 8 January 2001.
  31. DOUGLAS TUDHOPE, D. C. Semantically Indexed Hypermedia: Linking Information Disciplines. *ACM Computing Surveys* 31, 4es (December 1999).
  32. E. JAMES WHITEHEAD, J., AND WIGGINS, M. WEBDAV: IETF Standard for Collaborative Authoring on the Web. *IEEE Internet Computing* (September/October 1998), 34–40.
  33. EL-BELTAGY, S., DEROURE, D. C., AND HALL, W. A Multiagent system for Navigation Assistance and Information Finding. In *The Fourth International Conference on the Practical Application of Intelligent Agents and Multi-Agent Technology* (April 1999), pp. 281–295.
  34. ENGELBART, D. A Conceptual Framework for the Augmentation of Man's Intellect. *Vistas in Information Handling* 1, 9 (1963).
  35. ENGELBART, D. C. Knowledge-Domain Interoperability and an Open Hypermedia System. In *Proceedings of the Conference on Computer-Supported Collaborative Work* (Los Angeles, CA, October, 7-10, 1990), pp. 143–156.
  36. FERRAILOLO, J. Scalable Vector Graphics (SVG) 1.0 Specification. W3C Candidate Recommendations

- are available at <http://www.w3.org/TR>, 2 November 2000.
37. HALASZ, F. Reflections on NoteCards: Seven Issues for the Next Generation of Hypermedia Systems. *Communications of the ACM* 31, 7 (July 1988), 836–852.
  38. HALASZ, F., AND SCHWARZ, M. The Dexter Hypertext Reference Model. *Communications of the ACM* 37, 2 (February 1994), 30–39. Edited by K. Grønbaeck and R. Trigg.
  39. HARDMAN, L., BULTERMAN, D. C. A., AND ROSSUM, G. The Amsterdam Hypermedia Model: Adding Time and Context to the Dexter Model. *Communications of the ACM* 37, 2 (February 1994), 50–62.
  40. HARDMAN, L., VAN OSSENBRUGGEN, J., RUTLEDGE, L., MULLENDER, K. S., AND BULTERMAN, D. C. A. Do You Have the Time? Composition and Linking in Time-based Hypermedia. In *Proceedings of the 10th ACM conference on Hypertext and Hypermedia* [4], pp. 189–196. Edited by Klaus Tochtermann, Jorg Westbomke, Uffe K. Will and John J. Leggett.
  41. HILL, G., AND HALL, W. Extending the Microcosm Model to a Distributed Environment. In *Proceedings of the ACM European Conference on Hypermedia Technology (ECHT'94)* [2], pp. 32–40.
  42. HUNTER, J., AND LAGOZE, C. Combining RDF and XML Schemas to Enhance Interoperability Between Metadata Application Profiles. In *The Tenth International World Wide Web Conference* [43], pp. 457–466.
  43. IW3C2. *The Tenth International World Wide Web Conference* (Hong Kong, May 1-5, 2001).
  44. JORDAN, D. S., RUSSELL, D. M., JENSEN, A.-M. S., AND ROGERS, R. A. Facilitating the development of representations in hypertext with IDE. In *Hypertext '89 Proceedings* (Pittsburgh, Pennsylvania, November 5-8, 1989), ACM, ACM Press, pp. 93–104.
  45. KARVOUNARAKIS, G., CHRISTOPHIDES, V., PLEXOUSAKIS, D., AND ALEXAKI, S. Querying Community Web Portals. <http://www.ics.forth.gr/proj/isst/RDF/RQL/rql.html>.
  46. KLYNE, G. CC/PP Attribute Vocabularies. Work in progress. W3C Working Drafts are available at <http://www.w3.org/TR>, 21 July 2000.
  47. LANDOW, G. P. Relationally encoded links and the rhetoric of hypertext. In *Hypertext '87 Proceedings* [1], pp. 331–343.
  48. LUESEBRINK, M. C. The moment in hypertext: a brief lexicon of time. In *The Proceedings of the Ninth ACM Conference on Hypertext and Hypermedia* [3], pp. 106–112. Edited by Kaj Grønbaeck, Elli Mylonas and Frank M. Shipman III.
  49. MALCOLM, K. C., POLTROCK, S. E., AND SCHULER, D. Industrial Strength Hypermedia: Requirements for a Large Engineering Enterprise. In *Third ACM Conference on Hypertext Proceedings (Hypertext'91)* (San Antonio, Texas, December 15-18, 1991), ACM, ACM Press, pp. 13–24.
  50. MARSHALL, C., SHIPMAN, F., AND COOMBS, J. VIKI: Spatial Hypertext Supporting Emergent Structure. In *Proceedings of the ACM European Conference on Hypermedia Technology (ECHT'94)* [2], pp. 13–23.
  51. MARSHALL, C. C. Exploring representation problems using hypertext. In *Hypertext '87 Proceedings* [1], pp. 253–268.
  52. MILLER, J., RESNICK, P., AND SINGER, D. Rating Services and Rating Systems (and Their Machine Readable Descriptions), Version 1.1, October, 31, 1996. W3C Recommendations are available at <http://www.w3.org/TR>.
  53. NACK, F., AND HARDMAN, L. Denotative and Connotative Semantics in Hypermedia. Submitted to ACM Hypertext 2001, August 14-18, 2001.
  54. NANARD, J., AND NANARD, M. Should Anchors Be Typed Too — An Experiment with MacWeb. In *Fifth ACM Conference on Hypertext Proceedings (Hypertext'93)* (Seattle, Washington, November 14-18, 1993), ACM, ACM Press, pp. 51–62.
  55. NILSSON, M., HJELM, J., AND OHTO, H. Composite Capabilities/Preference Profiles: Requirements and Architecture. Work in progress. W3C Working Drafts are available at <http://www.w3.org/TR>, 21 July 2000.
  56. NÜRNBERG, P. J., LEGGETT, J. J., AND WILL, U. K. An Agenda for Open Hypermedia Research. In

- The Proceedings of the Ninth ACM Conference on Hypertext and Hypermedia* [3]. Edited by Kaj Grønbaeck, Elli Mylonas and Frank M. Shipman III.
57. RUTLEDGE, L., BAILEY, B., VAN OSSENBRUGGEN, J., HARDMAN, L., AND GEURTS, J. Generating Presentation Constraints from Rhetorical Structure. In *Proceedings of the 11th ACM conference on Hypertext and Hypermedia* [5], pp. 19–28.
  58. RUTLEDGE, L., VAN OSSENBRUGGEN, J., HARDMAN, L., AND BULTERMAN, D. C. Structural Distinctions Between Hypermedia Storage and Presentation. In *Proceedings of ACM Multimedia* (November 1998), ACM Press, pp. 145–150.
  59. SMITH, J. B., WEISS, S. F., AND FERGUSON, G. J. A hypertext writing environment and its cognitive basis. In *Hypertext '87 Proceedings* [1], pp. 195–214.
  60. TRIGG, R. *A Network-Based Approach to Text Handling for the Online Scientific Community*. PhD thesis, University of Maryland, November 1983. University of Maryland Technical Report, TR-1346.
  61. TRIGG, R., SUCHMAN, L., AND HALASZ, F. Supporting Collaboration in NoteCards. In *Proceedings of the Conference on Computer-Supported Collaborative Work* (Austin, Texas, December 3–5, 1986), pp. 147–153.
  62. VAN HARMELEN, F., AND HORROCKS, I. Reference description of the DAML+OIL ontology markup language. <http://www.daml.org/2000/12/reference.html>, December 2000. Contributors: Tim Berners-Lee, Dan Brickley, Dan Connolly, Mike Dean, Stefan Decker, Pat Hayes, Jeff Heflin, Jim Hendler, Deb McGuinness, Lynn Andrea Stein.
  63. VAN OSSENBRUGGEN, J. *Processing Structured Hypermedia — A Matter of Style*. PhD thesis, Vrije Universiteit, Amsterdam, The Netherlands, April 10, 2001. Also available on <http://www.cwi.nl/~jrvosse/thesis/>.
  64. VAN OSSENBRUGGEN, J., GEURTS, J., CORNELISSEN, F., RUTLEDGE, L., AND HARDMAN, L. Towards Second and Third Generation Web-Based Multimedia. In *The Tenth International World Wide Web Conference* [43], pp. 479–488. This is a revised version of CWI technical report INS-R0025.
  65. W3C. Synchronized Multimedia Integration Language (SMIL) 1.0 Specification. W3C Recommendations are available at <http://www.w3.org/TR/>, June 15, 1998. Edited by Philipp Hoschka.
  66. W3C. Resource Description Framework (RDF) Model and Syntax Specification. W3C Recommendations are available at <http://www.w3.org/TR/>, February, 22, 1999. Edited by Ora Lassila and Ralph R. Swick.
  67. W3C. Resource Description Framework (RDF) Schema Specification 1.0. W3C Candidate Recommendations are available at <http://www.w3.org/TR/>, 27 March 2000. Edited by Dan Brickley and R.V. Guha.
  68. W3C. XHTML 1.0: The Extensible HyperText Markup Language: A Reformulation of HTML 4.0 in XML 1.0. W3C Recommendations are available at <http://www.w3.org/TR/>, January 26, 2000.
  69. WANG, W., AND HAAKE, J. M. Flexible Coordination with Cooperative Hypermedia. In *The Proceedings of the Ninth ACM Conference on Hypertext and Hypermedia* [3], pp. 245–255. Edited by Kaj Grønbaeck, Elli Mylonas and Frank M. Shipman III.
  70. WHITEHEAD, E. J. Control choices and network effects in hypertext systems. In *Proceedings of the 10th ACM conference on Hypertext and Hypermedia* [4], pp. 75–82. Edited by Klaus Tochtermann, Jorg Westbomke, Uffe K. Will and John J. Leggett.