

Named Graphs, Provenance and Trust

Jeremy J. Carroll
Hewlett-Packard Labs
Bristol, UK
jjc@hpl.hp.com

Pat Hayes
IHMC, Florida, USA
phayes@ihmc.us

Christian Bizer
Freie Universität Berlin
Germany
chris@bizer.de

Patrick Stickler
Nokia, Finland
patrick.stickler@nokia.com

ABSTRACT

The Semantic Web consists of many RDF graphs nameable by URIs. This paper extends the syntax and semantics of RDF to cover such Named Graphs. This enables RDF statements that describe graphs, which is beneficial in many Semantic Web application areas. As a case study, we explore the application area of Semantic Web publishing: Named Graphs allow publishers to communicate assertional intent, and to sign their graphs; information consumers can evaluate specific graphs using task-specific trust policies, and act on information from those Named Graphs that they accept. Graphs are trusted depending on: their content; information about the graph; and the task the user is performing. The extension of RDF to Named Graphs provides a formally defined framework to be a foundation for the Semantic Web trust layer.

Categories and Subject Descriptors

I.2.4 [Artificial Intelligence]: Knowledge Representation Formalisms and Methods; K.6.5 [Management of Computing and Information Systems]: Security and Protection; H.3.3 [Information Storage and Retrieval]: Information Search and Retrieval—*selection process*

General Terms

Languages, Security

Keywords

RDF, Semantic Web, Provenance, Trust

1. INTRODUCTION

A simplified view of the Semantic Web is a collection of web retrievable RDF documents, each containing an RDF graph. The RDF Recommendation [4, 12, 26, 32], explains the meaning of any one graph, and how to merge a set of graphs into one, but does not provide suitable mechanisms for talking about graphs or relations between graphs. The ability to express meta-information about graphs is required for:

Data syndication systems need to keep track of provenance information, and provenance chains.

Restricting information usage Information providers might want to attach information about intellectual property rights or their

Copyright is held by the International World Wide Web Conference Committee (IW3C2). Distribution of these papers is limited to classroom use, and personal use by others.

WWW 2005, May 10-14, 2005, Chiba, Japan.
ACM 1-59593-046-9/05/0005.

privacy preferences to graphs in order to restrict the usage of published information [18, 34].

Access control A triple store may wish to allow fine-grain access control, which appears as metadata concerning the graphs in the store [28].

Signing RDF graphs As discussed in [13], it is necessary to keep the graph that has been signed distinct from the signature, and other metadata concerning the signing, which may be kept in a second graph.

Stating propositional attitudes such as modalities and beliefs [27].

Scoping assertions and logic where logical relationships between graphs have to be captured [6, 25, 31].

Ontology versioning and evolution OWL [19] provides various properties to express metadata about ontologies. In OWL Full, these ontologies are RDF graphs. Ontology evolution is discussed in [17].

RDF reification has well-known problems in addressing these use cases as previously discussed in [16]. To avoid these problems several authors propose quads [3, 20, 28, 33], consisting of an RDF triple and a further URIref or blank node or ID. The proposals vary widely in the semantic of the fourth element, using it to refer to information sources, to model IDs or statement IDs or more generally to ‘contexts’.

We propose a general and simple variation on RDF, called *named* RDF graphs. A Named Graph is an RDF graph which is assigned a name in the form of a URIref. The name of a graph may occur either in the graph itself, in other graphs, or not at all. Graphs may share URIrefs but not blank nodes.

Named Graphs can be seen as a reformulation of quads in which the fourth element’s distinct syntactic and semantic properties are clearly distinguished, and the relationship to RDF’s triples, abstract syntax and semantics is clearer.

As in both [37, 40] Named Graphs are treated as first class objects. The key contribution of this paper over and above such earlier work is the observation that the single feature, of graph naming, is the crucial one, but that the complex semantic theories of [37, 40] principally act as a barrier to deployment.

Named Graphs are a deliberately small step on top of the RDF and OWL Recommendations. This allows their use with tools built as implementing those recommendations, in a backward compatible way, with little or no code modifications.

The first half of the paper covers: the abstract and concrete syntaxes for Named Graphs; their semantics and the relationship with RDF and OWL; the relationship with TRIPLE [40] and with Guha and Fikes’ contexts [37]; and query languages for Named Graphs, including SPARQL [36].

The second half describes how Named Graphs can be used for Semantic Web publishing, looking in particular on provenance tracking and how it interacts with the choices consumers of Semantic Web information make about which information to trust. We provide a vocabulary for Semantic Web publishing with its formal semantics. The vocabulary includes support for digital signatures and addresses performative acts, such as asserting RDF.

2. ABSTRACT SYNTAX AND SEMANTICS

RDF syntax is based on a mathematical abstraction: an RDF graph is defined as a set of triples. These graphs are stored in documents which can be retrieved from URIs on the Web. Often these URIs are also used as a name for the graph, for example with an `owl:imports`. To avoid confusion between these two usages we distinguish between Named Graphs and the RDF graph that the Named Graph encodes or represents. A Named Graph is an entity with two functions *name* and *rdfgraph* defined on it which determine respectively its name, which is a URI, and the RDF graph that it encodes or represents. These functions assign a unique name and RDF graph to each Named Graph, but Named Graphs may have other properties.

More formally, let U be the set of all URIreferences, B an infinite set of RDF blank nodes, and L the set of all legal RDF literals (all three sets as defined in [32]); U , B and L are pairwise disjoint; let $V = U \cup B \cup L$ be the set of *nodes*; then the set $T = V \times U \times V$ is the set of all RDF triples.¹ The set of RDF graphs G is the power set of T . A Named Graph is a pair $ng = (n, g)$ with n in U and g in G . We write $name(ng) = n$ and $rdfgraph(ng) = g$. To enforce the blank node scoping rules ([26]) we make the global assumption that blank nodes cannot be shared between different Named Graphs, that is, if ng and ng' are different Named Graphs then the sets of blank nodes which occur in triples in $rdfgraph(ng)$ and in $rdfgraph(ng')$ are disjoint.

All of the above definitions may be relativized to a particular set of URIrefs, or to a particular set of Named Graphs. Any set of Named Graphs can be thought of as a partial function from U into the power set of T .

3. CONCRETE SYNTAXES

A concrete syntax for Named Graphs has to exhibit the name, the graph and the association between them. We offer three concrete syntaxes: TriX and RDF/XML both based on XML; and TriG as a compact plain text format.

The TriX [16] serialization is an XML format which corresponds fairly directly with the abstract syntax, allowing the effective use of generic XML tools such as XSLT, XQuery, while providing syntax extensibility using XSLT. TriX is defined with a short DTD, and also an XML Schema.

In this paper we use TriG as a compact and readable alternative to TriX. TriG is a variation of Turtle [5] which extends that notation by using '{' and '}' to group triples into multiple graphs, and to precede each by the name of that graph. The following TriG document contains two graphs. The first graph contains information about itself. The second graph refers to the first one (namespace prefix definitions omitted).

```
:G1 { _:Monica ex:name "Monica Murphy" .
      _:Monica ex:email
        <mailto:monica@murphy.org> .
      :G1 pr:disallowedUsage pr:Marketing }
```

¹We have removed the legacy constraint that a literal cannot be the subject of a triple.

```
:G2 { :G1 ex:author :Chris .
      :G1 ex:date "2003-09-03"^^xsd:date }
```

Named Graphs are backward compatible with RDF. A collection of RDF/XML [4] documents on the Web map naturally into the abstract syntax, by using the first `xml:base` declaration in the document or the URL from which an RDF/XML file is retrieved as a name for the graph given by the RDF/XML file. Using RDF/XML has disadvantages:

- The set of Named Graphs is in many documents rather than one.
- The known constraints and limitations of RDF/XML apply. For instance, it is not possible to serialize graphs with certain predicate URIs, nor is it possible to use literals as subjects.
- The URI at which an RDF/XML document is published is used for three different purposes: as a retrieval address, with an operational semantics; as a means of identifying the document; and as a means of identifying the graph described by the document. There is potential for confusion between these three uses.

None of these disadvantages is present in TriX and TriG. In balance, the major advantage of using RDF/XML is the deployed base, and current technology.

4. THE SEMANTICS OF NAMED GRAPHS

The semantics of graph naming are a simple semantic extension of the RDF(S) semantics: we will say that an RDF(S) interpretation I (as in [26]) *conforms* with a set of Named Graphs N when: For every Named Graph $ng \in N$, $name(ng)$ is in the vocabulary of I and $I(name(ng)) = ng$. Note that the Named Graph itself, rather than the RDF graph it intuitively 'names', is the denotation of the name. We consider the RDF graph to be related to the Named Graph in a way analogous to that in which a class extension is related to a class in RDFS. This intensional (c.f. [26]) style of modelling allows for distinctions between several copies of a single RDF graph (with distinct names) and avoids pitfalls arising from accidental identification of similar Named Graphs.

The RDF documentation [32] defines a notion of graph equivalence, which treats two RDF graphs which differ only in the identity of their blank nodes as being the 'same' graph. We will make a similar assumption, ignoring the mathematical details of 'renaming' functions; in practice, this amounts to permitting RDF processors to freely re-name any blank node identifiers when required in order to maintain the no-sharing condition.

The intuitive meaning of a Named Graph G is the standard RDF meaning [26] of its associated RDF graph $rdfgraph(G)$, which we will refer to as the *graph extension*. Any assertions in RDF about the graph structure of Named Graphs are understood to be referred to these graph extensions, just as the meanings of the RDFS class vocabulary are referred to relationships between the class extensions. As an example of this meaning, we can define two properties `rdfg:subGraphOf` and `rdfg:equivalentGraph`, with semantics defined as follows:

$$\langle f, g \rangle \text{ is in } IEXT(I(\text{rdfg:subGraphOf})) \\ \text{iff } rdfgraph(f) \text{ is a subset of } rdfgraph(g)$$

where the subset holds in a manner performing any necessary blank node renaming, as discussed above. Formally, the condition is that there is a renaming mapping m on the blank nodes of $rdfgraph(f)$ such that the RDF graph $m(rdfgraph(f))$ is a subset of $rdfgraph(g)$.

$$\langle f, g \rangle \text{ is in } IEXT(I(\text{rdfg:equivalentGraph})) \\ \text{iff } rdfgraph(f) = rdfgraph(g)$$

where, again, identity is understood as renaming blank nodes as appropriate.

We consider three further issues of detail in the relation between Named Graphs and RDF and OWL: the open world assumption; RDF reification, and OWL imports.

4.1 The Open World Assumption

Both RDF and OWL operate with the open world assumption. RDF Concepts [32] says:

RDF is an open-world framework that allows anyone to make statements about any resource. In general, it is not assumed that complete information about any resource is available.

The OWL Guide [41]:

OWL makes an open world assumption. That is, descriptions of resources are not confined to a single file or scope. While class C_1 may be defined originally in ontology O_1 , it can be extended in other ontologies.

As is clear from these quotations, openness here means that a description of a resource is considered to be open-ended. Actual web objects such as files and RDF graphs can however be identified and rigidly named, so that the name uniquely identifies the resource. Named Graphs utilize this ability to attach a name rigidly to a graph. Thus the mapping between names and graphs fixes the graph corresponding to a name in a rigid, non-extensible way. Two different Web documents asserting different graphs named by the same URI contradict one another. However, two different graphs with different names may make statements about the same resources. Thus the Named Graph framework facilitates the open world of the Semantic Web; not only can different people make different (hopefully complementary) statements about the same resource, but it is possible to keep these statements separate, and it is possible to combine them. The choice of which of these two is more appropriate is explicitly application specific.

Summarizing, if document A contains a graph g named u making statements about a resource r , a further document B that is consistent with A cannot use the name u for a different graph g' . However, B can contain a graph g' named u' making further statements about r . Thus the Named Graphs framework maintains the open-world framework of RDF, but treats graph naming as a form of rigid identification.

4.2 RDF Reification

A 'reified statement' [26] is a single RDF statement described and identified by a URIreference. Within the framework of this paper, it is natural to think of this as a Named Graph containing a single triple. With this convention, the subject of `rdg:subGraphOf` can be a reified triple, and the property can be used to assert that a Named Graph contains a particular triple. This provides a useful connection with the traditional use of reification and a potential migration path. However, the semantics of a single triple graph differ from the (lack of) semantics offered to a reified statement by the RDF recommendation [26], better addressing traditional uses of reification such as providing metadata about triples and quoting.

4.3 OWL Imports

OWL imports processing uses the URI object of an `owl:imports` triple to locate an additional RDF/XML file to be included in an ontology, as in [35], with K a collection of RDF graphs:

K is imports closed iff for every triple in any element of K of the form $x \text{ owl:imports } u$. then K contains a graph that is the result of the RDF processing

of the RDF/XML document, if any, accessible at u into an RDF graph.

Using the Named Graphs it is more natural to use the name of a graph as the object of `owl:imports`; so that the notion of imports closure is applied to a collection K of Named Graphs, and the definition is reworded as:

K is imports closed iff for every triple in any element of K of the form $x \text{ owl:imports } u$. then K contains a graph that is named u .

The URI u still may act as a locator, used to retrieve an RDF/XML document that is parsed to give a graph named u . The retrieval is unnecessary if the graph is available through other means, e.g., a cache (like with Jena's `OntDocumentManager`), or a local copy, or as part of a TriX document. There is a consistency question: do two different copies of a Named Graph agree? This can perhaps be resolved by phrasing it as: does a copy of a Named Graph agree with the graph found by the retrieval action?

5. RELATED WORK

Previous authors of research work addressing the semantics of a collection of documents on the Semantic Web have tended to have rich theories for addressing the relationship between multiple contexts.

5.1 TRIPLE

TRIPLE [40] provides graphs named with resources, and a Horn clause language for defining inferences etc., e.g.

```
@dfki:documents {
  dfki:d_01_01 [
    dc:title → "TRIPLE";
    dc:creator → "Michail Sintek";
    dc:creator → "Stefan Decker" ;
    ... ] .
  ∀ S,D search(S,D) ←
    D[dc:subject → S].
}
```

By mixing up the data representation (i.e., the Named Graph), the implementation of core RDF and DAML-OIL semantics (through Horn rules) and application semantics (through further Horn rules), TRIPLE becomes, as they say a 'novel query and transformation language'. Horn rules are allowed to reference data from multiple models. In as much as TRIPLE could be seen as mandating a single approach to implementing RDF(S) and OWL semantics (Horn rules), this must be seen as a weakness. The ongoing work on query languages for the Semantic Web indicates that other developers are more comfortable with a specification that does not presuppose a Horn implementation, but permits different developers to implement in different ways. Named Graphs can be seen as taking one aspect of this language, noting that it is particular useful, and not addressed by the Semantic Web recommendations, and pursuing that.

5.2 Contexts in RDF

Guha and Fikes [37] provide contexts, aggregate contexts, lifting rules, selective importing, preference rules, etc. They modify the RDF model theory to have additional context parameters both in the abstract syntax being interpreted and in the universe of interpretation. They interpret sets of graphs, rather than an individual graph. Unfortunately this step is sufficiently large to require significant new effort for implementors of RDF and OWL inference systems.

Their approach shares with ours the style of expressing some of the richer semantic constraints as extensions which constrain interpretations of certain new vocabulary (for them, e.g., `importsFrom`, for us e.g. `signature`).

A significant difference is the approach to aggregation. For Guha and Fikes certain contexts are aggregate contexts, which use lifting rules, possibly simple imports, possibly complex non-monotonic rules, to combine RDF data from multiple sources. They have some universal restrictions built into the built into the model theory, for example, lifting rules must be defined within their target aggregate context. For us, aggregation is only ever a monotonic merging operation, but the choice of what to aggregate is seen as a pragmatic application level decision.

We find their approach to be overly complex. Feigenbaum [22] suggests that for Semantic Web research that the “Path of maximal return is more knowledge not more logic”. Unlike Guha and Fikes we do not propose complex logic for contexts, merely the minimum step needed to record knowledge about provenance and other aspects of graphs needed for applications which need to address problems of trust. Using knowledge recorded with Named Graphs, applications will be able to use heuristics appropriate to them, to select the graphs they wish to trust for specific purposes.

The simple approach that we take permits substantially quicker deployment of applications that need to take provenance information into account, uses the flexibility and expressiveness of RDF, and is, we believe, fully adequate for Semantic Web applications in the near future. Web technology, designed to be deployed on a World Wide scale, needs to put a high value on simplicity, and on incremental steps. This ensures enough development effort can be made, in a number of systems, with different implementation strategies, to support the widespread deployment needed for a Web. The first steps of the Semantic Web are completed: systems implementing the RDF and OWL recommendations are deployed. Knowledge is published on the Semantic Web in these formats. To be effective, proposals for new Semantic Web features must build on these foundations, and must be parsimonious in the additional implementation effort required. A key feature of Named Graphs, lacking in TRIPLE or Guha’s contexts work, is parsimony.

5.3 RDF Dataset and SPARQL

A more recent development of the work in this paper is the RDF Dataset used in recent drafts of the SPARQL query language, defined in [36] as:

An RDF dataset is a set = $\{G, (u_1, G_1), (u_2, G_2), \dots, (u_n, G_n)\}$ where G and each G_i are graphs, and each u_i is a URI. Each u_i is distinct.

G is called the background graph. G_i are named graphs.

The main innovation over our work is the addition of the background graph. This provides backward compatibility with RDF without Named Graphs, and allows the Named Graphs functionality of SPARQL to be optional. Within the Named Graph framework presented here, the background graph of SPARQL could be implemented by using a distinguished name.

The addition of the background graph to a collection of Named Graphs may have the side effect of reintroducing some of the difficulties that Named Graphs address. For example, merging both background graphs and Named Graphs from different repositories, while maintaining provenance information, may prove difficult.

6. QUERY LANGUAGES

There are two current query languages for Named Graphs: RDFQ and TriQL.

RDFQ [42] uses an RDF vocabulary to structure queries. Queries can be constrained to Named Graphs matching one or more graph templates.

The following RDFQ query (serialized using Turtle [5]) identifies people having email addresses, selecting and extracting the person identifier and email address value pairs; furthermore, the query is restricted to statements occurring in graphs asserted by Chris after January 31, 2003:

```
[ :select ( "person" "email" );
  :graph [ex:author doc:Chris;
         ex:date
         [ :gt "2003-01-31"^^xsd:date ] ];
  :target [ :id "person";
          ex:email [ :id "email" ] ] ] .
```

TriQL [8] is a graph patterns based query language inspired by RDQL [39]. A graph pattern consists of a set of triple patterns and an optional graph name.

The following TriQL query has similar intent.

```
SELECT ?person ?email
WHERE ?graph ( ?person ex:email ?email )
( ?graph ex:author doc:Chris .
  ?graph ex:date ?date )
AND ?date > "2003-01-31"^^xsd:date
```

The example query uses two graph patterns. The variable `?graph` is bound to the names of all graphs that contain information about email addresses. The second pattern restricts `?graph` to graphs fulfilling both triple patterns.

The W3C is developing a new RDF query language SPARQL [36], which will also allow querying across multiple Named Graphs. RDFQ and TriQL predate SPARQL and we expect that SPARQL will supersede both languages once it has become a final W3C recommendation.

7. IMPLEMENTATIONS

Because Named Graphs are only a small addition on top of the Semantic Web recommendations it is easy to implement them using existing Semantic Web tools.

7.1 NG4J

One of these extensions is NG4J [10] which builds on the Jena Semantic Web toolkit [15]. NG4J provides developers with an API for manipulating sets of Named Graphs. It implements the TriX and the TriG syntax and the TriQL query language. A set of Named Graphs can also be viewed and manipulated as a provenance-enabled Jena model, allowing applications to track the origin of statements. By retrofitting Jena with an extended abstract syntax while staying compatible with the existing Jena API, NG4J aims at providing an migration path for existing applications based on Jena.

7.2 Jena MultiModel

One application which uses Named Graphs is a faceted browser [38], <http://www.swed.org.uk/>. This harvests RDF graphs from potentially many sites, and stores them in a `MultiModel` object which embodies the Named Graph abstraction on top of Jena’s `Model` class [15], which implements the RDF abstract syntax [32]. The source of any triple can be used during the faceted browse for visual styling of that part of the data. The end-user can apply varying levels of trust to different information presented on a single page. The style indicates the different authors, who can be treated with varying levels of caution.

8. SEMANTIC WEB PUBLISHING

One application area for Named Graphs is publishing information on the Semantic Web. This scenario implies two basic roles embodied by humans or their agents: Information providers and information consumers. Information providers publish information together with meta-information about its intended assertional status. Additionally, they might publish background information about themselves, e.g., their role in the application area. They may also decide to digitally sign the published information. Information providers have different levels of knowledge, and different intentions and different views of the world. Thus seen from the perspective of an information consumer, published graphs are claims by the information providers, rather than facts. An information consumer may accept some of these claims and reject others. We represent these choices by introducing the concept of the information consumer *accepting* Named Graphs.

Different tasks require different levels of trust. Thus information consumers will use different trust policies to decide which graphs should be accepted and used within the specific application. These trust policies depend on the application area, the subjective preferences and past experiences of the information consumer and the trust relevant information available. A naive information consumer might for example decide to trust all graphs which have been explicitly asserted. This trust policy will achieve a high recall rate but is easily undermineable by information providers publishing false information. A more cautious consumer might require graphs to be signed and the signers to be known, for example, through a Web-of-Trust mechanism. This policy is harder to undermine, but also likely to exclude relevant information, published by unknown information providers.

8.1 Accepting Graphs

Thus, a set of Named Graphs N has not been given a single formal meaning. Instead, the formal meaning depends on an additional set $A \subset \text{domain}(N)$. A identifies some of the graphs in the set as *accepted*. Thus there are $2^{|\text{domain}(N)|}$ different formal meanings associated with a set of Named Graphs, depending on the choice of A . The meaning of a set of accepted Named Graphs $\langle A, N \rangle$ is given by taking the graph merge $\bigcup_{a \in A} N(a)$, and then interpreting that graph with the RDF semantics [26], subject to the additional constraint that all interpretations I conform with N . Named Graphs can be used with any of the various levels of semantic theories for RDF: simple, RDF, RDFS or datatyped interpretations from [26], or OWL Full interpretations from [35]. It is a deliberate choice to work in this way with the deployed Semantic Web Recommendations, rather than inventing a new semantics with special features, perhaps from modal logic, to reflect potential conflict between different graphs on the Semantic Web.

The choice of A reflects that the individual graphs in the set may have been provided by different people, and that the information consumers who use the Named Graphs make different choices as to which graphs to believe. Thus we do not provide one correct way to determine the ‘correct’ choice of A , but provide a vocabulary for information providers to express their intentions, and suggest techniques with which information consumers might come to their own choice of which graphs to accept. Issues as to how to resolve conflicts between different graphs, and how to determine A , are seen as pragmatic issues, to be dealt with by application developers, rather than logical issues to be dealt with by formal semantics. A motivation is that different applications will have different tolerances to errors, inconsistencies and variability between the data, and a unified formal approach is likely to be overkill for some, yet may miss key features required by another (e.g., some more formal

approaches to context [37, 40, 23] fail to address digital signatures, vital for financially sensitive applications).

8.2 Authorities, Authorization and Warrants

Information providers using RDF do not have any explicit way to express any intention concerning the truth-value of the information described in a graph; RDF does not provide for the expression of *propositional attitudes*, such as asserting, denying, commenting on, or otherwise expressing an opinion about the content of a graph. Information consumers may require this, however. Note that this is in addition to trust policies, and may be required in order to put such policies into operation. For example, a simple policy could be: believe anything asserted by a trusted source. In order to apply this, it is necessary to have a clear record of what is *asserted* by the source. Not all information provided by a source need be asserted by that source. We propose here a vocabulary and a set of concepts designed to enable the uniform expression of such propositional attitudes using named graphs.

We take three basic ideas as primitive: that of an *authority*, a relationship of *authorizing*, and a *warrant*. An authority is a ‘legal person’; that is, any legal or social entity which can perform acts and undertake obligations. Examples include adult humans, corporations and governments. The ‘authorizing’ relationship holds between an authority or authorities and a Named Graph, and means that the authority in some sense commits itself to the content expressed in the graph. Whether or not this relationship in fact holds may depend on many factors and may be detected in several ways (such as the Named Graph being published or digitally signed by the authority). Finally, a warrant is a resource which records a particular propositional stance or intention of an authority towards a graph. A warrant asserts (or denies or quotes) a Named Graph and is authorized by an authority. One can think of warrants as a way of reducing the multitude of possible relationships between authorities and graphs to a single one of authorization, and also as a way of separating questions of propositional attitude from issues of checking and recording authorizations. The separation of authority from intention also allows a single warrant to refer to several graphs, and for a warrant to record other properties such as publication or expiry date.

To describe the two aspects of a warrant we require vocabulary items: a property `swp:authority` (where `swp:` is a namespace for Semantic Web publishing) relating warrants to authorities, and another to describe the attitude of the authority to the graph being represented by the warrant. We will consider two such intentions expressed by the properties `swp:assertedBy` and `swp:quotedBy`. These take a named graph as a subject and a `swp:Warrant` as object; `swp:authority` takes a warrant as a subject and a `swp:Authority` as an object. Each warrant must have a unique authority, so `swp:authority` is an OWL functional property. Intuitively, `swp:assertedBy` means that the warrant records an endorsement or assertion that the graph is true, while `swp:quotedBy` means that the graph is being presented without any comment being made on its truth. This latter is particularly useful when republishing graphs as part of a syndication process, the original publisher may assert a news article, but the syndicator, acting as a common carrier, merely provides the graph as they found it, without making any commitment as to its truth. Warrants may also be signed, and the property `swp:signatureMethod` can be used to identify the signature technique.

8.3 Warrant Descriptions as Performatives

A warrant, as described above, is a social act. However, it is often useful to embody social acts with some record; for example a con-




Figure 1: The Semantic Web Publishing Vocabulary

tract (which is a social act) may be embodied in a document, which is identified with that act, and is often signed. In this section, we introduce the notion of a *warrant graph*, which is a Named Graph describing a warrant, that is identified with the social act. Thus, this is a resource which is both a `swp:Warrant` and an `rdf:Graph`. Consider a graph containing a description of a warrant of another Named Graph, such as:

```
{ :G2 swp:assertedBy _:w .
  _:w rdf:type swp:Warrant2 .
  _:w swp:authority _:a .
  _:a rdf:type swp:Authority .
  _:a foaf:mbox <mailto:chris@bizer.de> }
```

The graph is true when there is a genuine warrant; but so far we have no way to know whether this is in fact the case. A slight modification identifies the graph with the warrant itself:

```
:G1 { :G2 swp:assertedBy :G1 .
      :G1 swp:authority _:a .
      _:a foaf:mbox <mailto:chris@bizer.de> }
```

Suppose further that such a *warrant graph* is, in fact, authorized by the authority it describes - in this case, by Chris Bizer, the owner of the mailbox: this might be established for example by being published on Chris' website, or by being digitally signed by him, or in some other way, but all that we require here is that it is in fact true. Under these circumstances, the warrant graph has the intuitive force of a first-person statement to the effect "I assert :G2" made by Chris.

In natural language, the utterance of such a self-describing act is called a *performative*; that is, an act which is performed by saying that one is doing it. Other examples of performatives include promising, naming and, in some cultures, marrying [2]. The key point about performatives are that while they are descriptions of themselves, they are not only descriptions: rather, the act of uttering the performative is understood to be the act that it describes. Our central proposal for how to express propositional attitudes on the Web is to treat a warrant graph as a record of a performative act, in just this way.³ With this convention, Chris can assert the graph :G2 by authorizing the warrant graph shown above, for by doing so he creates a warrant: the warrant graph becomes the (self-describing) warrant of the assertion of :G2 by Chris. In order for others to detect and confirm the truth of this warrant requires some way to check or confirm the relationship of authorization, of course: but the qualification of the warrant graph as a warrant depends only on the relationship holding.

A graph describing a warrant is not required to be self-describing in order to be true (it may be true by virtue of some other warrant)

²The type triples are implied by domain and range constraints and can be omitted.

³The Bank of England uses this technique, by having each twenty pound note bear the text: "I promise to pay the bearer on demand the sum of twenty pounds."

and a warrant graph may not in fact be a performative warrant (if it is not authorized by the authority it claims). In the latter case the graph must be false, so self-describing warrant graphs whose authorization cannot be checked should be treated with caution. The warrant graph may itself be the graph asserted. Any Named Graph which has a warrant graph as a subgraph and is appropriately authorized satisfies the conditions for being a performative warrant of itself. For example:

```
:G2 { :Monica ex:name "Monica Murphy" .
      :G2 swp:assertedBy :G2 .
      :G2 swp:authority _:a .
      _:a foaf:mbox
      <mailto:patrick.stickler@nokia.com> . }
```

when authorized by Patrick Stickler, becomes a performative warrant for its own assertion, as well as being warranted by the earlier example. As this example indicates, a Named Graph may have a number of independent warrants.

8.4 Publishing with Signatures

Information providers may decide to digitally sign graphs, when they wish to allow information consumers to have greater confidence in the information published. For instance, if Patrick has an X.509 certificate [30], he can sign two graphs in this way:

```
:G1 { :Monica ex:name "Monica Murphy" .
      :G1 swp:assertedBy _:w1 .
      _:w1 swp:authority _:a .
      _:a foaf:mbox
      <mailto:chris@bizer.de> }
:G2 { :G1 swp:quotedBy :G2 .
      :G1 swp:digestMethod
      swp:JjcRdfC14N-sha1 .
      :G1 swp:digest
      "...^^xsd:base64Binary .
      :G2 swp:assertedBy :G2 .
      :G2 swp:signatureMethod
      swp:JjcRdfC14N-rsa-sha1 .
      :G2 swp:signature
      "...^^xsd:base64Binary .
      :G2 swp:authority _:s .
      _:s foaf:mbox
      <mailto:patrick.stickler@nokia.com> .
      _:s swp:certificate
      "...^^xsd:base64Binary }
```

Note that :G2 is a warrant graph. The `swp:signature` gives a binary signature of the graph related to the warrant⁴. The canonicalization algorithms and the signature method which have been used to calculate the signature are indicated by the value of the `swp:signatureMethod` property. SWP uses a similar mechanism as XML-Signature [21] for signing several resources using a single signature: Including the graph digest of :G1 into :G2 and signing :G2 afterwards also assures the integrity of :G1.

The information publisher indicates the methods used for forming digests and signatures. We require the methods to be identified by literal URIs, which can be dereferenced on the Web to retrieve a document, describing the method in detail. The signature method `swp:JjcRdfC14N-rsa-sha1`, for example, specifies a variation of the graph canonicalization algorithms provided in [13], and chooses one of the digest/signature method combinations defined by XML-Signature [21]. Rather than make a set of decisions about digest and signature methods, SWP provides terms for describing the chosen combination.

The publisher may choose to sign graphs to ensure that the max-

⁴It is necessary to exclude the last `swp:signature` triple, from the graph before signing it: this step needs to be included in the method.

imum number of Semantic Web agents believe them and act on the publication. Using signatures does not modify the theoretical semantics of assertion, which is boolean; but it will modify the operational semantics, in that without signatures, any assertions made, will only be acted on by the more trusting Semantic Web information consumers, who do not need verifiable information concerning who is making them.

The formal semantics of the Semantic Web publishing vocabulary are described in Section 9.

8.5 The Information Consumer

The information consumer needs to decide which graphs to accept. This decision may depend on information concerning who said what, and whether it is possible to verify such information. It may also depend on the content of what has been said. We consider the use case in which an information consumer has read a set of Named Graphs off the Web. In terms of the semantics of Named Graphs (Section 8.1), the information consumer needs to determine the set A . Information about the graphs may be embedded within the set of Named Graphs, hence most plausible trust policies require that we are able to provisionally understand the Named Graphs in order to determine, from their content, whether or not we wish to accept them. This is similar to reading a book, and believing it either because it says things you already believe, or because the author is someone you believe to be an authority: either of these steps require reading at least some of the book.

The trust policy an information consumer chooses for determining his set of accepted graphs depends on the application area, his subjective preferences and past experiences and the trust relevant information available. Trust policies can be based on the following types of information [11]:

First-hand information published by the actual information provider together with a graph, e.g., information about the intended assertional status of the graph or about the role of the information provider in the application domain. Example policies using the information provider's role are: "Prefer product descriptions published by the manufacturer over descriptions published by a vendor" or "Distrust everything a vendor says about its competitor."

Information published by third parties about the graph (e.g., further assertions) or about the information provider (e.g., ratings about his trustworthiness within a specific application domain). Most trust architectures proposed for the Semantic Web fall into this category [1, 7, 24]. These approaches assume explicit and domain-specific trust ratings. Providing such ratings and keeping them up-to-date puts an unrealistically heavy burden on information consumers in many application domains.

The content of a graph together with rules, axioms and related content from graphs published by other information providers. Example policies following this approach are "Believe information which has been stated by at least 5 independent sources." or "Distrust product prices that are more than 50% below the average price."

Information created in the information gathering process like the retrieval date and retrieval URL of a graph or whether a warrant attached to a graph is verifiable or not.

Example trust policies are found in [9, 11].

We sketch an algorithm that allows the agent to implement a trust policy of trusting any RDF that is explicitly asserted. This is intended to be illustrative, in the sense that different agents should have different trust policies, and these will need different algo-

rithms. More cautious variation may require performative assertions or digital signatures.

The agent has an RDF knowledge base, K , which may or may not be initially populated. The agent is presented with a set of Named Graphs N , and augments the knowledge base with some of those graphs (determining the set A of accepted graphs).

1. Set $A := \phi$
2. Non-deterministically choose $n \in \text{domain}(N) - A$, if no choices remain terminate.
3. Set $K' := K \cup N(n)$, provisionally assuming $N(n)$.
4. If K' entails:
 $n \text{ swp:assertedBy } _ : w$.
then set $K := K'$ and $A := A \cup \{n\}$, otherwise backtrack to 2.
5. Repeat from 2.

If initially K is empty, then the first graph added to K will be one that includes its own assertion, by an arbitrary warrant and authority. All such graphs will be added to K , as will any that are asserted as a consequence of the resulting K . The algorithm is equivalent to one that seeks to accept a graph by finding a statement of its assertion either within itself, or within some other accepted graph, or the initial knowledge base. The algorithm is sound with respect to the goal of only adding graphs that are explicitly asserted, as verified by step 4. However, it is incomplete against the same criterion, since two graphs each of which explicitly assert the other, would satisfy the criterion if both were accepted, but the algorithm misses that. We see the self-asserting performative warrant as the basic communicative act, and more sophisticated phrasings (such as the mutually asserting graphs), are less likely to be understood.

At step 4, $_ : w$ is unconstrained, reflecting the simple policy of trust everybody. A slightly more sophisticated query could implement a policy that, for example, only trusted a set of named individuals, or require that any self-asserting graph actually be a warrant graph.

This algorithm does not take consistency into account. As we merge internally consistent graphs in step 3 we may introduce inconsistencies that occur between the graphs. In some cases, it may not be possible to even detect this, for example in OWL Full which has an undecidable theory. For a semantics with a complete and terminating consistency checker [14] (such as for OWL Lite), inconsistency could be detected immediately. We do not propose any particular response to inconsistency. Some applications, such as the faceted browser of [38], may not care, whereas others, may wish to use inconsistency to reject some of the graphs, in favour of a maximal consistent subset. Mechanisms such as those used in truth maintenance systems would be useful for these applications.

8.5.1 Using a Public Key Infrastructure

The trust algorithm above would believe fraudulent claims of assertion. That is, any of the Named Graphs may suggest that anyone asserted any of the graphs, whether or not that is true, and the above algorithm has no means of detecting that.

We have described how a publisher can sign their graphs and include such signatures in the published graphs. We will continue to explore the X.509 certified case; in general the PGP [44] case is similar, and the approach taken does not assume a particular PKI.

The earlier example can be checked by modifying the query in step 4 to be:

```
SELECT ?certificate ?method ?sign
WHERE ( ?w1 swp:assertedBy ?w1 .
        ?w1 swp:authority ?s .
        ?w1 swp:signatureMethod ?method .
        ?w1 swp:signature ?sign )
        ( ?s swp:certificate ?certificate )
```

where this is understood as being over the interpretation of the graph, rather than as a syntactic query over the graph. The signatures must be verified following the given method. If this verification fails then the graph is false and can be rejected. If the verification succeeds then the certification chain should be considered by the information consumer. If the agent trusts anyone in the certificate chain⁵, then the graph is accepted, otherwise not. More sophisticated algorithms would consider whether the person asserting the graph, who has now been verified, is in the group of persons which the information consumer trusts on the topic the graph discusses.

A graph may have more than one warrant. If any warrant contains an incorrect signature, then the warrant is simply wrong, and indicates data or algorithmic corruption. A graph containing such a warrant (but not necessarily the named graph misasserted) should be rejected. The choice of which warrant to check is nondeterministic and hence should consider any valid warrant whose certification chain is trusted.

9. FORMAL SEMANTICS OF PUBLISHING AND SIGNING

This section provides an extension of RDF semantics [26] which: allows persons to be members of the domain of discourse; allows interpretations to be constrained by the identifying information in a digital certificate; allows the `swp:assertedBy` triple to have a *performative* semantics; and makes `swp:signature` triples true or false depending on whether the signature is valid or not. Together these extensions underpin the publishing framework of the previous section.

9.1 Persons in the Domain of Discourse

The two frameworks of digital signatures we have considered both tie a certificate to a legal person (e.g., a human or a company), or similar. In X.509, a certificate includes a distinguished name [29, 43], which is chosen to adequately identify a legal person, and is verified as accurate by the certification authority. In PGP, a certificate contains unspecified identifying information, “such as his or her name, user ID, photograph, and so on” [44]; this is usually an e-mail address.

The class extension of `swp:Authority` is constrained to be a set P of legal persons and software agents acting on behalf of legal persons. Thus, our formal semantics requires the universe of discourse to contain such persons as resources. Such a requirement goes beyond the usual ‘logical’ bounds of model-theoretic semantics. We expect that Web languages will further extend their semantics into the real world of agents, acts and things as they become applied in real-world applications. This first step is, in itself, not very interesting since we have not constrained which person in the real world corresponds to which URIref or blank node in the graph.

The second step is to constrain the property extension of `swp:-certificate` to $\{(p, c) | p \in P, c \text{ a sequence of binary octets, with } c \text{ being an X.509 or PGP certificate for } p\}$. The binary octets can be represented in a graph using `xsd:base64Binary`. The interpretation of these sequences as X.509 is specified in [30], which gives a distinguished name from RFC 2253 [43], which identifies a person. If c gives a PGP certificate then given the potential vagueness of the identifying information we allow all pairs of in which the person matches the identifying information. For example, if the

⁵For PGP, the specific method of determining whether the certificate is trusted is different.

identifying information is only a GIF image, then all people who look like that image are paired with the certificate.⁶

This definition does *not* depend on whether or not the certificate is trusted. If the graph containing the `swp:certificate` triple is accepted, using mechanisms such as those discussed in Section 8.5, then the triple’s meaning is as above. The certificate chain in the certificate is only checked when deciding which graphs to accept.

9.2 Performative warrants

A formal model-theoretic semantics specifies truth conditions. To fully capture the meaning of a performative, however, we need to go beyond truth-conditions, since the very same form of words may be true whoever uses them, but will only count as a performative if used by the authority it mentions. For example “Patrick promises...” uttered by Patrick is a promise - a performative act - but uttered by Christian is merely a description of the act; yet it may well still be true, and for the same reasons. We will deal with this by considering a warrant graph to be a ‘genuine’ warrant just when it describes its authority accurately, and to be true in any interpretation under which a genuine warrant actually exists.

The relationship of authorizing, and sets of authorities and warrants (from Section 8.2), are taken as primitive, and we will identify them, respectively, with the property extension of `swp:authority` and the class extensions of `swp:Authority` and `swp:Warrant`. All the remaining semantic conditions are defined in terms of these, so their correctness in any application depends on that of the interpretation of `swp:authority` together with its range and domain. Thus a triple

`ex:a swp:authority ex:b` .

is true in I just when $I(\text{ex:a})$ is a warrant which is authorized by $I(\text{ex:b})$.

The performative role of a properly authorized warrant graph can then be described by simply declaring that any Named Graph ng containing a triple

`name(ng) swp:authority bbb` .

is a warrant. Then any interpretation I of `rdggraph(ng)` (conforming to the naming of ng) under which ng is authorized by $I(\text{bbb})$ makes this triple true, and hence requires ng to be in $ICEXT(I(\text{swp:Warrant}))$: call this an *authorizing interpretation* of the Named Graph. Fixing the referent of the object of such a triple to be an authorizing authority thus means that the graph can be satisfied only by authorizing interpretations under which the Named Graph is a warrant.

The self-realizing quality of performatives is extended to the triples which express propositional attitudes by making these trivially self-fulfilling when they occur under the right conditions, in an authorized warrant graph. For example if ng is a warrant graph which contains a triple

`aaa swp:assertedBy bbb` .

where $I(\text{bbb}) = ng$, then if I is an authorizing interpretation of ng , then I must satisfy that triple; similarly for `swp:quotedBy` and indeed for any other property expressing a propositional attitude of an authority towards a graph.

Note that this does not imply that a Named Graph is *true* in an authorizing interpretation of a warrant which asserts it. The fact of an authority asserting a graph can be true independently of the actual truth of the graph. However, the attitude expressed can be utilized by trust policies. It may be appropriate to treat graphs asserted by trusted authorities as being true, but not to extend this to

⁶This shows why it is unwise to only provide an image in your PGP certificate.

graphs quoted by trusted authorities. One could express this trust policy by a semantic rule to the effect that if I satisfies

```
aaa swp:assertedBy bbb .
bbb swp:authority ccc .
```

and $I(ccc)$ is trusted, then I satisfies $rd\text{fgraph}(I(aaa))$.

The algorithm for choosing which graphs to accept, presented in Section 8.5, interacts with this performative semantics, by essentially assuming that a graph has been asserted, and then verifying that in that case the performative is true.

Using `rdfs:subPropertyOf` or `owl:equivalentProperty` to introduce aliases of `swp:assertedBy` may be misleading and should be avoided. Information consumers should be suspicious of any graphs that attempt this, except when they are also asserted by the persons using the aliases so introduced.

9.3 Graph Digests and Signatures

The final specialized vocabulary we consider is that for graph digests and signatures. Strictly speaking this is not necessary for Semantic Web publishing, but just as a signed document has greater social force than an unsigned one, a signed `swp:assertedBy` triple is more credible than an unsigned one. Thus, this section is specifically intended to be used to sign graphs that are either the subject of, or include `swp:assertedBy` triples.

The two properties `swp:digest` and `swp:signature` are treated in a similar fashion: we start with the simpler `swp:digest`.

A pair (g, d) is in the property extension of `swp:digest`, if and only if,

1. d is a finite sequence of octets.
2. There is a pair (g, m) in the property extension of `swp:digestMethod`, and m is a URI which can be dereferenced to get a document.
3. The method described in the document retrieved from m calculates the digest d for the graph $I(g)$.

This means that an `swp:digest` triple is true if and only if the value is the appropriate digest. Hence, if the graph which is the subject of the triple has been tampered with, such tampering is detected by the `swp:digest` triple being false.

Similarly, a pair (w, s) is in the property extension of `swp:signature`, if and only if,

1. s is a finite sequence of octets.
2. There is a pair (w, m) in the property extension of `swp:signatureMethod`, and m is a URI which can be dereferenced to get a document.
3. There is a pair (w, a) in the property extension of `swp:authority` and a pair (a, c) in the property extension of `swp:certificate`, and c is a finite sequence of octets.
4. There is a pair (g, w) in the property extension of `swp:quotedBy` or `swp:assertedBy`, and $I(g)$ is a Named Graph.
5. The method described in the document retrieved from m calculates the signature s for the graph $I(g)$ using c understood as an X.509 or PGP certificate.

This definition does not depend upon verifying the certificate chain for c .

10. CONCLUSIONS

Having a clearly defined abstract syntax and formal semantics Named Graphs provide greater precision and potential interoperability than the variety of *ad hoc* RDF extensions currently used. Combined with specific further vocabulary, this will be beneficial in a wide range of application areas and will allow the usage of a common software infrastructure spanning these areas.

The ability of self-reference combined with the Semantic Web Publishing vocabulary addresses the problem of differentiating asserted and non-asserted forms of RDF and allows information providers to express different degrees of commitment towards published information.

Linking information to authorities and optionally assuring these links with digital signatures gives information consumers the basis for using different task-specific trust-policies. We have shown how operational trust can depend on what is being said, rather than simply depending on who said it, and the trust-rating of the author.

Named Graphs provide a high-value but small and incremental change to the Semantic Web Recommendations. Thus they should be preferred over more complex, all-embracing approaches to context that are more likely to face substantial barriers to adoption.

Further related work can be found at the TriX and Named Graphs web-site <http://www.w3.org/2004/03/trix/>.

11. ACKNOWLEDGEMENTS

Thanks to the W3C Semantic Web Interest Group for their interest and comments on this work as it has developed. During the development of this work, Christian Bizer has been a visiting researcher at HP Labs in Bristol and Jeremy Carroll has been a visiting researcher at ISTI, CNR in Pisa. Christian Bizer thanks his host Andy Seaborne. His visit was supported by the Leonardo Da Vinci grant no. D/2002/EX-22020. Jeremy Carroll thanks his host Oreste Signore. Pat Hayes' work was supported in part by the US Department of Defense under contract 2507-282-22. Thanks to anonymous reviewers of various versions of this paper for their helpful feedback.

12. REFERENCES

- [1] R. Agrawal, P. Domingos, and M. Richardson. Trust Management for the Semantic Web. In *Proceedings of the 2nd ISWC*, 2003.
- [2] J. L. Austin. *How to do things with words*. Harvard University Press, 1962.
- [3] D. Beckett. Redland Notes - Contexts. <http://www.redland.opensource.ac.uk/notes/contexts.html>, 2003.
- [4] D. Beckett. RDF/XML Syntax Specification (Revised). <http://www.w3.org/TR/rdf-syntax-grammar/>, 2004.
- [5] D. Beckett. Turtle - Terse RDF Triple Language. <http://www.ilrt.bris.ac.uk/discovery/2004/01/turtle/>, 2004.
- [6] T. Berners-Lee. Notation 3. <http://www.w3.org/DesignIssues/Notation3>, 1998.
- [7] C. Bizer. Semantic Web Trust and Security Resource Guide. <http://www.wiwiss.fu-berlin.de/suhl/bizer/SWTSGuide>, 2004.
- [8] C. Bizer. TriQL - A Query Language for Named Graphs. <http://www.wiwiss.fu-berlin.de/suhl/bizer/TriQL/>, 2004.
- [9] C. Bizer. TriQLP - A Query Language for Querying Named Graphs Published by Untrustworthy Sources. <http://www.wiwiss.fu-berlin.de/suhl/bizer/triqlp/>, 2004.
- [10] C. Bizer and R. Cyganiak. Ng4j - named graphs api for jena. <http://www.wiwiss.fu-berlin.de/suhl/bizer/ng4j/>, 2005.

- [11] C. Bizer and R. Oldakowski. Using Context- and Content-Based Trust Policies on the Semantic Web. In *13th World Wide Web Conference, WWW2004 (Poster)*, 2004.
- [12] D. Brickley and R. V. Guha. RDF Vocabulary Description Language 1.0. <http://www.w3.org/TR/rdf-schema/>, 2004.
- [13] J. J. Carroll. Signing RDF Graphs. In *2nd ISWC*, volume 2870 of *LNCS*. Springer, 2003.
- [14] J. J. Carroll and J. De Roo. Web Ontology Language (OWL) Test Cases. <http://www.w3.org/TR/owl-test/>, 2004.
- [15] J. J. Carroll, I. Dickinson, C. Dollin, D. Reynolds, A. Seaborne, and K. Wilkinson. Jena: Implementing the Semantic Web Recommendations. In *13th World Wide Web Conference, WWW2004*, 2004.
- [16] J. J. Carroll and P. Stickler. TriX: RDF Triples in XML. Technical Report HPL-2003-268, HP Labs, 2004.
- [17] C. Catton and D. Shotton. The use of Named Graphs to enable ontology evolution. <http://www.bioimage.org/pub/paradigm.htm>, 2004.
- [18] Creative Commons Website. <http://creativecommons.org/>, 2003.
- [19] M. Dean and G. Schreiber. OWL Web Ontology Language Reference. <http://www.w3.org/TR/owl-ref/>, 2004.
- [20] E. Dumbill. Tracking Provenance of RDF Data. Technical report, ISO/IEC, 2003.
- [21] D. Eastlake, J. Reagle, and D. Solo. XML-Signature Syntax and Processing, RFC 3275. <http://www.w3.org/TR/xmlsig-core/>, 2002.
- [22] E. Feigenbaum. Keynote. In *International Semantic Web Conference*, 2004.
- [23] A. Gangemi and P. Mika. Understanding the semantic web through descriptions and situations. In *Proceedings of DOA/CoopIS/ODBASE 2003 Confederated International Conferences DOA, CoopIS and ODBASE*, LNCS. Springer, 2003.
- [24] J. Golbeck, B. Parsia, and J. Hendler. Trust Networks on the Semantic Web. In *Proceedings of the 7th International Workshop on Cooperative Intelligent Agents, CIA2003*, 2003.
- [25] R. V. Guha. *Contexts: A Formalization and Some Applications*. PhD thesis, Stanford, 1995.
- [26] P. Hayes. RDF Semantics. <http://www.w3.org/TR/rdf-mt/>, 2004.
- [27] A. Ibrahim. Agent Communication Languages (ACL). <http://www.engr.uconn.edu/~ibrahim/publications/acl.htm>, 2000.
- [28] Intellidimension. RDF Gateway - Database Fundamentals. <http://www.intellidimension.com/pages/rdfgateway/dev-guide/db/db.rsp>, 2003.
- [29] ITU-T. The Directory – Models X.501, 1993.
- [30] ITU-T. Information Technology - Open Systems Interconnection - The Directory Authentication Framework. X.509, 1997.
- [31] G. Klyne. Circumstance, provenance and partial knowledge. <http://www.ninebynine.org/RDFNotes/UsingContextsWithRDF.html>, 2002.
- [32] G. Klyne and J. J. Carroll. Resource Description Framework (RDF): Concepts and Abstract Syntax. <http://www.w3.org/TR/rdf-concepts/>, 2004.
- [33] R. MacGregor and I.-Y. Ko. Representing Contextualized Data using Semantic Web Tools. In *Practical and Scalable Semantic Systems (workshop at 2nd ISWC)*, 2003.
- [34] M. Marchiori. The platform for privacy preferences. <http://www.w3.org/TR/P3P/>, 2002.
- [35] P. F. Patel-Schneider, P. Hayes, and I. Horrocks. OWL Web Ontology Language Semantics and Abstract Syntax. <http://www.w3.org/TR/owl-semantics/>, 2004.
- [36] E. Prud'hommeaux and A. Seaborne. SPARQL Query Language for RDF. <http://www.w3.org/TR/2005/WD-rdf-sparql-query-20050217/>, Feb 2005.
- [37] R. M. R. Guha and R. Fikes. Contexts for the semantic web. In *Proceedings of ISWC'2004*, 2004.
- [38] D. Reynolds, P. Shabajee, S. Cayzer, and D. Steer. Semantic portals demonstrator - lessons learnt. <http://www.w3.org/2001/sw/Europe/reports/demo.2.report/>, 2004.
- [39] A. Seaborne. RDQL - A Query Language for RDF. <http://www.w3.org/Submission/2004/SUBM-RDQL-20040109>, 2004.
- [40] M. Sintek and S. Decker. Triple - a query, inference, and transformation language for the semantic web. In *Proceedings of ISWC'2002*, pages 364–378, 2002.
- [41] M. K. Smith, C. Welty, and D. L. McGuinness. OWL Web Ontology Language Guide. <http://www.w3.org/TR/owl-guide/>, 2004.
- [42] P. Stickler. RDFQ. <http://sw.nokia.com/rdfq/RDFQ.html>, 2004.
- [43] M. Wahl, S. Kille, and T. Howes. Lightweight Directory Access Protocol (v3): UTF-8 String Representation of Distinguished Names, RFC 2253, 1997.
- [44] P. Zimmerman and Network Associates Inc. An Introduction to Cryptography. <ftp://ftp.pgpi.org/pub/pgp/6.5/docs/english/IntroToCrypto.pdf>, 1999.