

On the Integration of Topic Maps and RDF Data

Martin S. Lacher
Stanford University Database Group

Stefan Decker
Stanford University Database Group

Keywords: Topic maps; RDF; Queries; Interoperability; Contrasting approaches: Topic maps and RDF; Tool/implementation: RDF adapter for Topic maps; Experiments and demos: Topic map of World Fact Book and Open Directory RDF; semi-structured data; TM4J Topic map Engine; World Fact Book

Abstract

Topic Maps and RDF are two independently developed paradigms and standards for the representation, interchange, and exploitation of data or metadata about information resources. Each paradigm has established its own user communities. Each standard describes a graph-based data model with nodes and labeled arcs and one or more XML- or SGML-based serialization syntaxes. However, the two data models have significant conceptual differences. A central goal of both paradigms is to define an interchangeable format for the exchange of different kinds of data in a machine processable way on the Web. In order to prevent a partition of the Web into collections of incompatible resources, it is reasonable to seek ways for integration of Topic Maps with RDF. A first step is made by representing Topic Map information as RDF information and thus allowing Topic Map information to be queried by an RDF-aware infrastructure. To achieve this goal, the Topic Map graph model will be mapped to the RDF graph model. In order to stay as close to the original graph model as possible, the set of built-in node and arc types in Topic Maps is defined for RDF with an RDF Schema. The result of the mapping is an RDF-based internal representation of Topic Maps data that can be queried as an RDF source by an RDF-aware query processor.

On the Integration of Topic Maps and RDF Data

§ 1 Introduction

Different Communities are currently working on the vision of a Semantic Web: the idea of having data on the Web defined and linked in a way that it can be used by machines not just for display purposes, but for automation, integration and reuse of data across various applications. In order to make this vision a reality for the Web, supporting standards, technologies and policies must be designed to enable machines to make more sense of the Web, with the result of making the Web more useful for humans. One issue for the Semantic Web is how to allow for interoperable representations of data on the Web. TopicMaps[ISO 13520] and RDF[Lassila and Swick, 1999] are two independently developed standards, which can be used to represent data on the web in an interoperable fashion. Both standards have established a large user community and will most likely be building blocks of the future Semantic Web. To prevent a partition of the Semantic Web into incompatible subsets, ways for interoperation of overlapping standards like RDF and Topic Maps have to be found. By Interoperability, we mean for example that any Topic Map source of data can be queried with an RDF-aware query infrastructure and vice versa. Both directions are equally important, as both standards have their advantages and disadvantages and are equally likely to be used on the future Semantic Web. We chose to begin with the approach of making Topic Map sources queriable for an RDF infrastructure, because the RDF community has established a query infrastructure (eg.[Decker et. al., 1998]), which can be reused for querying Topic Map resources. The Topic Map community is in the process of standardizing a query language¹. Other approaches that make RDF sources available to Topic Map aware query infrastructure have been proposed and their relation to this work is presented in [Section 5 \[below\]](#).

Our approach to integration of Topic Maps and RDF data uses the layered approach to data interoperability proposed in [Melnik and Decker, 2000]. This approach splits data models into different layers, much like the layers in a network protocol stack. This layered model is useful for understanding complex data model interoperation, since the integration problem is broken into layers. An introduction to the layered approach to data interoperability is given in [Section 2 \[below\]](#). We make a Topic Map RDF-queriable by performing a mapping between the two data models on the object layer, which can in both cases be seen as a graph. Thus in fact, our mapping is a mapping between two types of graphs. The mapping is performed by modeling the Topic Map graph with an RDF graph. On top of the graph layer, there may be additional semantics, which we do not consider in this paper. For example, the graph may be used to represent UML data, DAML+OIL data or Topic Map data. Figure 1 shows an overview of the architecture that we have in mind for the integration of different sources.

our integration approach in more detail including a small example mapping. [Section 4 \[below\]](#) describes the implementation of our approach. [Section 5 \[below\]](#) gives a brief overview of related work. [Section 6 \[below\]](#) presents the application example of a common query to a Topic Map and the Open Directory. Finally, in [Section 7 \[below\]](#) we summarize our contributions.

§ 2 Overview of the data models

In this section, we will give a brief overview of the RDF and Topic Map data models with respect to the layered model introduced in [Melnik and Decker, 2000].

2.1 *The layered interoperability model*

The layered model of data interoperability in [Melnik and Decker, 2000] breaks up the problem of data model integration into a stack of layers which are quasi-independent from each other. This approach resembles the ISO protocol stack for network interoperation. The different layers presented are from bottom to top, the syntax layer, the object layer and the semantic layer. Each of those layers actually has sublayers, but we do not require such a detailed perspective on the layers here. The syntax layer is concerned with a serialization syntax for persistent storage of data. The object layer is concerned with how to assign identity to objects or how binary relations are represented. The semantic layer is concerned with the interpretation of the objects and their relationships.

We will not present details on each of the layers and their involvement in the mapping. The important essence is, that our approach works by performing a bijective graph transformation on the object layer, which can be performed quasi-independently from the other layers. This independence is possible, because any semi-structured data model [Suciu, 1998] can be represented as a directed graph, which is also the data model of RDF. Thus, any kind of semi-structured data model can be represented by RDF on the object layer. How the RDF graph is interpreted on a higher level can differ again for different data models. In this paper we will not consider the issue of mapping those higher level semantics. We will only look at RDF as the common denominator for data representation and query purposes.

2.2 *RDF*

The Resource Description Framework Model and Syntax Specification [Lassila and Swick, 1999], which became a World Wide Web Consortium (W3C) Recommendation in February 1999, defines the RDF data model and an XML-based serialization syntax. The RDF Data model is essentially a directed, labeled graph: it consists of entities, represented by unique identifiers, and binary relationships between those entities. In RDF, a binary relationship between two specific entities plus the entities itself is called a statement (or triple). Represented graphically, the source of the relationship is called the subject, the labeled arc is the predicate (also called property), and the relationship's destination is called the object of the statement. The RDF data model distinguishes between resources, which have URI identifiers, and literals, which are just strings. The subject and the predicate of a statement are always resources, while the object can be a resource or a literal.

Taking the perspective of the layered interoperability model, RDF has several possible syntaxes on the syntax layer, among which there is one XML syntax defined in [Lassila and Swick, 1999]. On the object layer, the RDF model is a directed graph, as described above.

2.3 *TopicMaps*

Topic Maps [ISO 13520] have also been standardized in 1999. A Topic Map is defined as

a collection of Topic Map documents, which adhere to a certain SGML syntax defined in the standard document. The SGML Syntax of those documents is described in the standard along with an informative conceptual model for memory representation of Topic Maps. To make Topic Maps applicable on the Web, the XML Topic Maps standard has been drafted [Pepper and Moore, 2001]. XTM defines an XML syntax for Topic Maps and gives a more specific data model of a Topic Map. Both the SGML syntax and the XML syntax incorporate syntax shortcuts for complex data model constructs. Moreover, Topic Maps define a multidimensional Topic space with interrelations between topics. The serialization and deserialization of a Topic Map is thus not straightforward, which is why guidelines for implementors of Topic Map software have been published in the form of processing models in [Biezunski and Newcomb, 2001] and [Biezunski and Newcomb, 2001a]. These processing models are very important since they are the only source of a valid mapping from the XTM syntax to a valid internal Topic Map representation. The processing model for XTM describes a graph based data model for Topic Maps. The graph model incorporates four different kinds of arcs as well as three different kinds of nodes. Possible syntaxes for Topic Map serialization are the SGML syntax defined in [ISO 13520] and the XML syntax first defined in [Pepper and Moore, 2001], which has been included in [ISO 13520] after publication. The data model of Topic Maps on the object layer is an undirected graph with certain types of arcs and nodes, as explained above. Moreover, arcs in a Topic Map graph can have arcs attached to them as well and the arcs have two distinct ends, which both have fixed labels according to the respective arc type.

§ 3 Integration Approach

Our general approach is that of modeling Topic Maps with the means and vocabulary that RDF gives us. This is an approach that has been termed "modeling the model" in [Moore, 2001]. The advantage of this approach is that the mapping preserves all information. In contrast, a semantic mapping could possibly lead to a loss of information in the mapping process.

3.1 *Semi-structured data*

Our integration goal is to generate a memory internal representation of a Topic Map, which can be queried with an RDF query infrastructure. This means that the surface syntax of the two data models is not of interest for our task. Thus, our approach is applicable for both the SGML syntax as well as the XML syntax. However, our implementation only considers the XML (XTM) syntax. We implemented the processing model proposed in [Biezunski and Newcomb, 2001a] to construct a Topic Map graph model from an XTM document.

RDF is closely related to the concept of semi-structured data, identified in the database community [Hammer et al, 1997] [Suciu, 1998] as a means for data integration [Garcia-Molina, 1995] [Papakonstantinou, 1995] and transformation [Abiteboul, 1997]. Any kind of data that can be represented as a graph is called semi-structured data. Thus, if heterogeneous data sources are transformed into a graph representation in some standard representation format, all this data can be queried with the same query infrastructure in the same query. This makes joint queries over multiple data sources possible.

RDF can be used to represent semi-structured data as a graph. This also applies to Topic Maps data, since there is a graph representation defined for Topic Maps [Biezunski and Newcomb, 2001a]. Topic Maps have the expressive power of a schema language and can be used to represent ontologies. An RDF adapter for Topic Maps makes a Topic Map information source RDF queryable.

We will now describe the different aspects of the representation of Topic Maps as RDF with respect to the layered data model described in [Melnik and Decker, 2000]

3.2 Object layer

The representation of Topic Maps as RDF is a graph transformation on the object layer. The object layer refers to characteristics such as how object identity is established or how binary relationships are established. These can be translated to graph characteristics of a graph representation of semi-structured data. The graph representations of RDF and Topic Maps are different in a number of characteristics. The following list describes how these differences are overcome.

Arc types: the Topic Map graph model presented in [Biezunski and Newcomb, 2001a] knows four different types of arcs. In RDF, there is a small set of built-in arcs (properties). The set of arcs can be expanded virtually without limit through the use of namespaces. We make use of the namespace capability of RDF to define the four types of arcs available in Topic Maps. The arc types are defined through an RDF Schema, shown in Figure 2.

Node types: the Topic Map graph model knows three different kinds of nodes. In RDF, there is a small set of built-in node types. Again, the use of namespaces allows the virtually limitless expansion of the node-type space. We make use of the namespace capability of RDF to define the three types of nodes available in Topic Maps. The node type definition is also shown in Figure 2 [below].

Object identity: in a Topic Map, a topic is uniquely identifiable through its basename and a namespace characterization. In RDF, each subject node in the graph has to be uniquely identifiable through a unique URI. The problem is, that in a Topic Map, the basename property is not mandatory. However, in the XML syntax, topics may have an additional ID for unique identification. So, for the unique identification of a node in an RDF graph representation of a Topic Map we chose to use the ID attribute, if it is available. If it is not available, we generate an ID.

Arc direction: in a Topic Map graph, arcs are undirected, but each arc end is labeled with a fixed label according to the type of arc. This resembles an implicit directionality, but arcs can still be traversed in either direction. RDF only allows directed arcs. With a given type of RDF arc, the direction is explicitly given and the end labels of the Topic Map arc are implicitly given. Keeping two directed arcs instead of one undirected arc would lead to consistency problems. Thus, we only keep one directed arc. This has to be considered when a query is formulated, as arcs have to be traversed in both directions then. The transition from undirected arcs to directed arcs is not a lossy transformation, since the arcs in the Topic Map graph are implicitly directed. The direction can be uniquely derived from the nodes which are attached to each arc.

Arcs and properties: in a Topic Map graph, arcs can have properties, i.e. outgoing arcs, as well. This is the case for the role label of an association member arc in a Topic Map graph. In the RDF graph we represent this by reifying the statement that includes the respective arc and the two adjacent nodes. An additional arc ending in another node is assigned to that reified statement.

3.3 Semantic layer

RDF can be the basis for an ontology definition language and Topic Maps can be seen as an ontology definition language. RDF requires additional vocabulary such as DAML+OIL for ontology definition and RDF itself merely provides the object layer in this data model stack. Topic Maps on the other side have richer semantics, and provide a number of features of an ontology definition language. For a comparison on the semantic layer, DAML+OIL based on RDF is a more appropriate candidate for a comparison with Topic Maps. However, this will not be investigated in this paper.

3.4 Integration Example

We will now present a small example for the representation of Topic Map data as an RDF graph. As a first preparatory step for our integration approach, we defined an RDF Schema which defines the node and arc types of a Topic Map graph. Figure 2 [below] shows the RDF schema definition.

```

<rdf:RDF
  xmlns:rdf="http://www.w3c.org/1999/02/22-rdf-syntax-ns#"
    xmlns:rdfs="http://www.w3c.org/2000/01/rdf-schema#"
  xmlns:tms="http://www.stanford.edu/rdf/tmmapping/tm-schema#"
  xmlns="http://www-db.stanford.edu/rdf/tmmapping/tm-schema#"
  >
  <rdfs:Class ID="topic" rdfs:comment="The class of topic
nodes"/>
  <rdfs:Class ID="association" rdfs:comment="The class of
association nodes"/>
  <rdfs:Class ID="scope" rdfs:comment="The class of scope
nodes"/>
  <rdf:Property ID="associationMember" rdfs:comment="The
association member arc">
    <rdfs:domain rdf:resource="#association"/>
    <rdfs:range>
      <rdf:Alt>
        <rdf:li rdf:resource="#association"/>
        <rdf:li rdf:resource="#topic"/>
      </rdf:Alt>
    </rdfs:range>
  </rdf:Property>
  <rdf:Property ID="associationScope" rdfs:comment="The
association scope arc">
    <rdfs:domain rdf:resource="#association"/>
    <rdfs:range rdf:resource="#scope"/>
  </rdf:Property>
  <rdf:Property ID="associationTemplate" rdfs:comment="The
association template arc">
    <rdfs:domain rdf:resource="#association"/>
    <rdfs:range rdf:resource="#topic"/>
  </rdf:Property>
  <rdf:Property ID="scopeComponent" rdfs:comment="The scope
component arc">
    <rdfs:domain rdf:resource="#scope"/>
    <rdfs:range>
      <rdf:Alt>
        <rdf:li rdf:resource="#association"/>
        <rdf:li rdf:resource="#topic"/>
      </rdf:Alt>
    </rdfs:range>
  </rdf:Property>
  <rdf:Property ID="roleLabel" rdfs:comment="The
association role label arc">
    <rdfs:range rdf:resource="#topic"/>
  </rdf:Property>
</rdf:RDF>

```

Figure 2: RDF Schema for an RDF-based Topic Map

For the actual construction of an RDF representation of a Topic Map graph, the next step is the generation of a graph representation from a (XTM) Topic Map document. For this purpose we implemented an API for Topic Maps which exposes a graph-based data structure and allows us to directly operate on the Topic Map constructs for the graph construction. The API also conforms with the processing model presented in [Biezunski and Newcomb, 2001a], which is required to generate a valid Topic Map graph from an abbreviated syntax. Figure 3 [below] shows a short snippet of a Topic Map with information from the CIA World Fact Book in the form of an XTM document. Processing this XTM document results in the graph shown in Figure 4 [below].

```

<topic id="denmark">
  <basename>
    <baseNameString>Denmark</baseNameString>
  </basename>
</topic>
<association id="denmark-has-petroleum">
  <member>
    <roleSpec>
      <topicRef xlink:href="#country"/>
    </roleSpec>
    <topicRef xlink:href="#denmark"/>
  </member>
  <member>
    <roleSpec>
      <topicRef xlink:href="#natural-resource">
    </roleSpec>
    <topicRef xlink:href="petroleum">
  </member>
</association>
<topic id="country"/>
<topic id="natural-resource"/>

```

Figure 3: XTM document subpart

After processing the XTM document snippet according to the processing model, the generated graph for this short XTM document snippet looks like this:

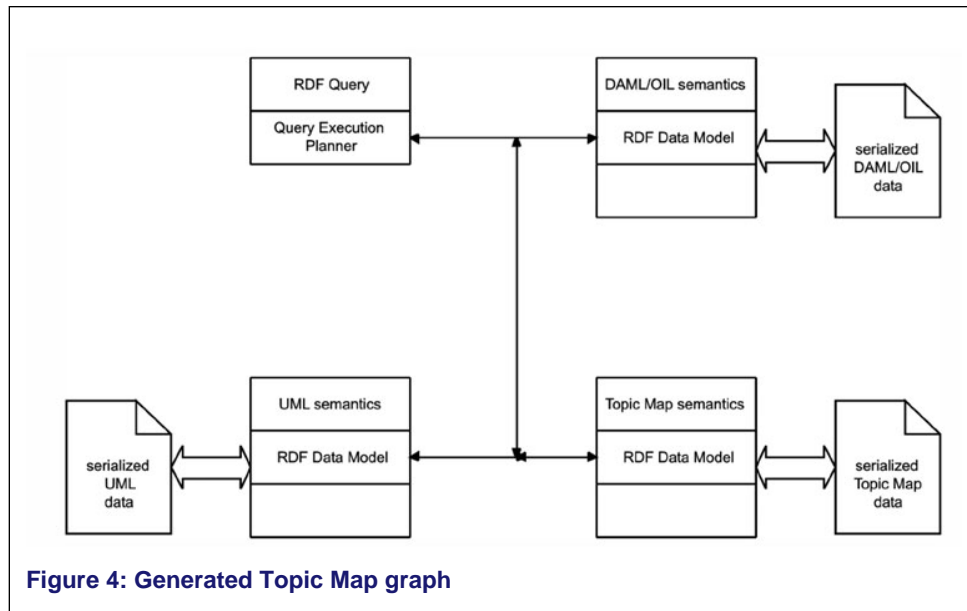


Figure 4 [above] shows the Topic Map graph that is generated according to the XTM processing model. The ellipses represent nodes, the lines represent arcs with different types. The role labels for association member arcs are connected to the arcs via another arc, the role label arc. The graph that is induced by the XTM snippet above basically represents a topic node that represents the subject Denmark. The graph also represents the fact that Denmark has petroleum as a natural resource. It also shows that the basename "Denmark" has been assigned to the Denmark topic.

We will now represent this graph as an RDF graph. In fact, the transformation of the graph is performed during the construction of the Topic Map graph according to the transformation guidelines presented above. To construct the graph, we generate RDF triples. Figure 5 [below] shows the mapped RDF graph.

2001a]. The construction of the graph model is performed through a graph-based API proposed in [Ahmed, 2001]. The implementation of this API simplifies the realization of the processing model, since the underlying data model is the same for both. Along with the creation of the API objects, an equivalent set of RDF triples is generated.

For parsing the XTM document we use a parser implemented in the TM4J Topic Map engine ². The SAX-based parser feeds events to our implementation of the processing model, which then constructs the RDF graph.

§ 5 Related Work

RDF and TopicMaps integration has been discussed for a while. In [Moore, 2001] two general approaches to the integration have been proposed. The first approach shows how Topic Maps can be modelled with RDF vocabulary and vice versa. The second approach shows how a semantic mapping between the two standards can be performed. Semantic mappings bear the disadvantage that inherently, the transformation is lossy and the transformation is not bijective.

Also, representing RDF data as Topic Map data is possible, but for the purpose of querying various sources through one query infrastructure, the inverse direction is the easier solution. RDF has the simpler data model, allowing more efficient and simpler storage and query facilities than TopicMaps. Pure syntax transformations have been proposed ³, but this approach disregards the need for a processing model to generate the Topic Map graph from the serialized syntax.

We have shown that from the point of view of an integrated Semantic Web it is desirable to be able to query a Topic Map source with an RDF query. This can be achieved if the Topic Map source itself represents its data as RDF data. The problem of integration of RDF and Topic Maps has been approached with little success so far. Most Integration approaches have lead to the conclusion that RDF is not expressive enough to represent Topic Maps. What we aim to achieve is not to convert a Topic Map document into a number of serialized RDF statements, which would render the document difficult to read. Instead we aim to generate an internal representation of a Topic Map, which is really a set of RDF statements. This way, a data source which stores Topic Map data can be queried as if it was an RDF source. Thus, what we need to achieve is a mapping of an internal Topic Map representation to an internal representation of a set of RDF statements.

§ 6 An Application Example: Combining a TopicMap with the Open Directory

As an example for the usefulness of our integration approach, consider the following scenario: We would like to find Web pages about travel in countries, which exploit petroleum as a natural resource. The available resources include a Topic Map constructed from the CIA world fact book ⁴, which includes general resources about countries, but no Web pages about travel. To retrieve the requested travel pages, we access the Open Directory collection of Web pages. The Open Directory is a large Web page directory constructed in a collaborative way by a large number of expert volunteers. The directory structure of the Open Directory is represented in RDF. With our integration approach, a query processor can now query both information sources and integrate the results into one query result. The distributed and heterogeneous nature of the information sources remains transparent to the user.

We will now show how the above query to the two information resources looks like. Figure 6 [below] shows an example of a query in F-Logic syntax, as introduced in [Decker et. al., 1998]. The query processing engine also proposed there, can answer the

following query:

```
FORALL pages <- Country, DMOZCountry Y,X, Z
Y[tms:roleLabel->country;rdf:object->Country]@CIA_WORLD_FACTBOOK
and
  X[tms:roleLabel->natural-resource;
   rdf:object->petroleum;
rdf:subject->Z[tms:associationMember->Country]@CIA_WORLD_FACTBOOK]
  @CIA_WORLD_FACTBOOK and
  Country[mapsTo->DMOZCountry] and
  DMOZCountry[Travel_and_Tourism
->dmozpage[links->pages]]@DMOZ.
```

Figure 6: Query in F-Logic syntax over DMOZ and RDF-based Topic Map

The query answers queries over two different sources: the CIA World Factbook and the DMOZ Open Directory. The structure of the query language mimics RDF and is `subject[predicate->object]@source`. The first part of the query retrieves all countries, which have petroleum as a natural resource. This part of the query can be answered from the CIA World Factbook Topic Map, in the RDF representation given above. We assume the existence of a name mapping, which resolves the naming differences between resources (`mapsTo` property). Now we are able to query the DMOZ data for travel information on this country. The result of the query is a list of DMOZ categories like "Top/Regional/Europe/Austria/Travel", "Top/Regional/Europe/France/Travel", etc. Please note that the query in figure 6 is simplified. A real working query will additionally have to deal with naming conventions of DMOZ and construct the "Travel_and_Tourisms" property from the DMOZCountry URI.

A graphical interface like the one presented in [Staab et. al., 2000] can simplify the query formulation for the user. It can also be ensured in the client query software that the query is broken up into subqueries which are directed to the right information sources.

§ 7 Conclusion

Interoperability is of greatest importance on the Semantic Web. We suggested a way to achieve interoperability between Topic Maps and RDF, which enables the joint querying of RDF and Topic Maps information sources. We achieved this by adopting an internal graph representation for Topic Maps, according to one of the processing models for Topic Maps that have been published. We perform a graph transformation to generate an RDF graph from the Topic Map graph representation. The Topic Map source can now be queried with an RDF query language together with RDF information sources. We see this as a first step towards the integration of the many heterogeneous information sources available on the Web today and in the future.

Notes

1. See <http://k42.empolis.co.uk/tmql.html>
 2. See <http://www.techquila.com/tm4j/>
 3. <http://lists.w3.org/Archives/Public/www-rdf-interest/2001Mar/0062.html>
 4. <http://www.cia.gov/cia/publications/factbook/>
-

Bibliography

- [**Abiteboul, 1997**] Serge Abiteboul, Sophie Cluet, Tova Milo: Correspondence and Translation for Heterogeneous Data. ICDT 1997: 351-363.
- [**Ahmed, 2001**] Khalil Ahmed: Developing a Topic Map Programming Model. In Proceedings of Knowledge Technologies 2001.
- [**Biezunski and Newcomb, 2001**] Michel Biezunski and Steven R. Newcomb: XML Topic Maps (XTM) Processing Model 1.0, Topicmaps.org Specification, 2001. See <http://www.topicmaps.org/xtm/1.0/xtmp1.html>
- [**Biezunski and Newcomb, 2001a**] Michel Biezunski and Steven R. Newcomb: Topicmaps.net's Processing Model for XTM 1.0, version 1.0.1, Topicmaps.net Specification, 2001. See <http://www.topicmaps.net/pmtm4.htm>
- [**Decker et. al., 1998**] Stefan Decker, Dan Brickley, Janne Saarela, Jürgen Angele: A Query and Inference Service for RDF. In Proceedings of the Query Languages Workshop '98.
- [**Garcia-Molina, 1995**] H. Garcia-Molina, J. Hammer, K. Ireland, Y. Papakonstantinou, J. Ullman, J. Widom: Integrating and Accessing Heterogeneous Information Sources in TSIMMIS. Proceedings of AAAI Spring Symposium on Information Gathering, 1995.
- [**Hammer et al, 1997**] J. Hammer, J. McHugh, H. Garcia-Molina: Semistructured Data: The TSIMMIS Experience. In Proceedings of the First East-European Workshop on Advances in Databases and Information Systems-ADBIS '97, St. Petersburg, Russia, September 1997
- [**ISO 13520**] Michel Biezunski, Steven Newcomb (Editors): ISO/IEC 13250 Topic Maps. 1999. See <http://www.y12.doe.gov/sgml/sc34/document/0129.pdf>
- [**Lassila and Swick, 1999**] Resource Description Framework (RDF) Model and Syntax Specification, <http://www.w3.org/TR/REC-rdf-syntax/> , 1999
- [**Melnik and Decker, 2000**] Sergey Melnik, Stefan Decker: A Layered Approach to Information Modeling and Interoperability on the Web. In Proceedings of the Workshop "ECDL 2000 Workshop on the Semantic Web", 2000.
- [**Moore, 2001**] Graham D. Moore: RDF and Topic Maps: An exercise in convergence. In Proc. of XML Europe 2001, Berlin, Germany, 2001.
- [**Papakonstantinou, 1995**] Y. Papakonstantinou, H. Garcia-Molina, J. Widom: Object Exchange Across Heterogeneous Information Sources, ICDE '95.
- [**Pepper and Moore, 2001**] Steve Pepper and Graham D. Moore (Editors): XML Topic Maps (XTM) 1.0, Topicmaps.org Specification, 2001. See <http://www.topicmaps.org/xtm/1.0/>
- [**Staab et. al., 2000**] Steffen Staab, J. Angele, Stefan Decker, Michael Erdmann, Andreas Hotho, Alexander Mädche, Hans-Peter Schnurr, Rudi Studer, York Sure: Semantic Community Web Portals. In: WWW9/Computer Networks (Special Issue: WWW9 - Proceedings of the 9th International World Wide Web Conference, Amsterdam, The Netherlands, May, 15-19, 2000
- [**Suciu, 1998**] Dan Suciu: An Overview of Semistructured Data. Published in SIGACT News, vol. 29, no. 4, pp. 28-38, December, 1998.
- [**XTMP**] The XTM 1.0 Processing Model, <http://www.topicmaps.org/xtm/1.0/xtmp1.html>

Extreme Markup Languages 2001

Montréal, Québec, August 14-17, 2001

This paper produced from XML source via XSL, Saxon and Apache FOP

Mulberry Technologies, Inc., August 2001