On the Integration of Topic Maps and RDF Data

Martin S. Lacher and Stefan Decker
Database Group
Stanford University
Stanford, CA 94306
{lacher,stefan}@db.stanford.edu

Abstract. Topic Maps and RDF are two independently developed paradigms and standards for the representation, interchange, and exploitation of model-based data on the web. Each paradigm has established its own user communities. Each of the standards allows data to be represented as a graph with nodes and labeled arcs which can be serialized in one or more XML- or SGML-based 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 knowledge 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, we map a Topic Map graph model to the RDF graph model. All information from the Topic Map is preserved, such that the mapping is reversible. The mapping is performed by modeling the graph features of a Topic Map graph model with an RDF graph. 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.

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 inter-operable representations of data on the Web. RDF [1] and Topic Maps [2] are two independently developed standards, which can be used to represent data on the web in an inter-operable 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 inter-operation 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. The reason for this is partly because the RDF community has established a query infrastructure (eg. [9]), which can be reused for querying Topic Map resources. The Topic Map community is in the process of standardizing a query language, commmercial packages already offer proprietary query languages [1]. 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 6.

Our approach to integration of Topic Maps and RDF data fllows the layered approach to data interoperability proposed in [12]. 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 complexity is broken into smaller problem parts. An introduction to the layered approach to data interoperability is given in Section 2. We make a Topic Map RDF-queriable by performing a mapping between the two data models on a layer, on which in both of the models, data is represented 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.

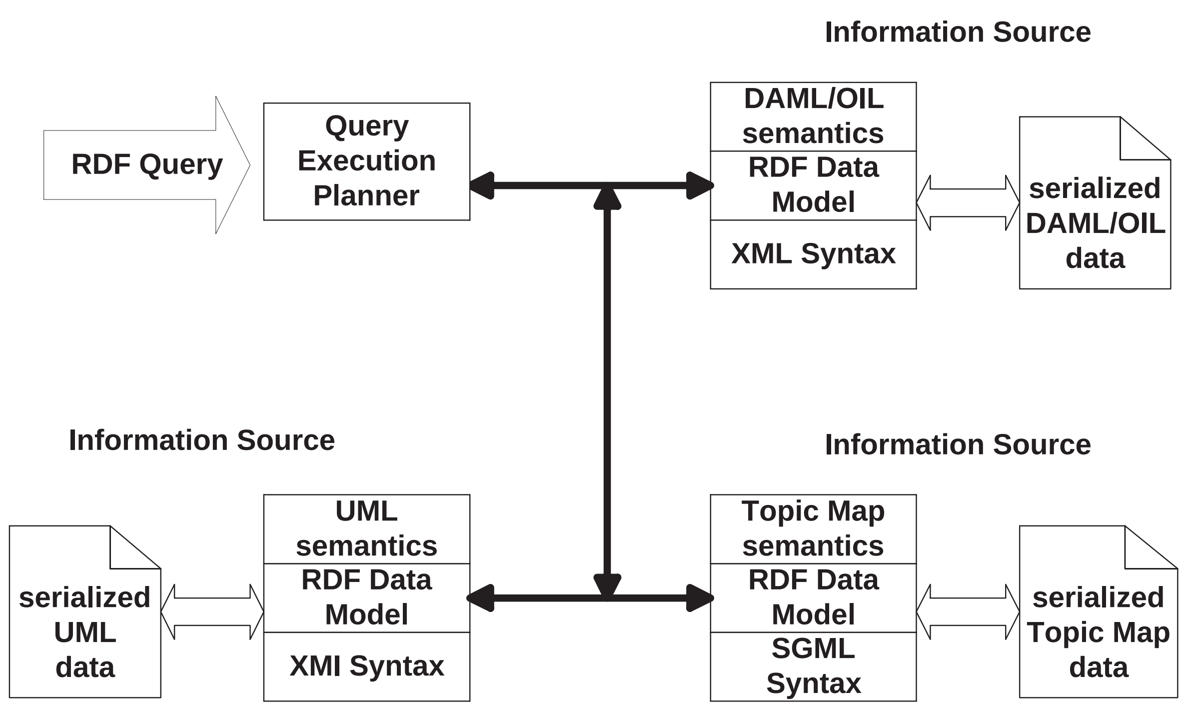

Figure 1: Overview of the integration of different data sources.

Each of the data sources in Figure 1 stores persistent data according to a certain serialization syntax. From each of these persistent data, a memory data model based on RDF as a low-level object model can be built. This RDF model in all information resources can then be queried by

[1]See http://k42.empolis.co.uk/tmql.html

an RDF-aware query infrastructure. This way, information sources with different model-based data representations can be integrated.

The remainder of this paper is organized as follows: We will first introduce the data models of RDF and Topic Maps with respect to the layered interoperability approach. General familiarity with RDF and Topic Maps is assumed. Thereafter, in Section 3, we will present our integration approach in more detail including two small exemplary mappings. Section 4 presents a real world application example for the joint querying of a Topic Map information source and an RDF information source. Section 5 shortly describes the implementation of our mapping approach. Section 6 gives a brief overview of related work. Finally, in Section 7 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 [12].

2.1 The layered interoperability model

The layered model of data interoperability in [12] 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 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 [15] 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. The Topic Map semantic on a higher level will thus be conserved with our mapping and only the representation on the object layer will be mapped to RDF.

2.2 RDF

The Resource Description Framework Model and Syntax Specification [1], which became a World Wide Web Consortium (W3C) Recommendation in February 1999, defines the RDF data model and a basic serialization syntax. The RDF Data model is essentially a directed, labeled graph: it consists of entities, identified by unique identifiers, and binary relationships between those entities. In RDF, a binary relationship between two specific entities is represented by a statement (or triple). An RDF statement can be represented in a graph as two nodes and a directed arc between the nodes. The node with the outgoing arc is called subject of the statement, the arc is called property and the node with the incoming arc is called 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 basic and on abbreviated syntax defined in [1]. On the object layer, the RDF model is a directed graph, as described above. The semantic layer of RDF is minimal. Together with additional languages like DAML+OIL, more more complex semantics can be expressed with RDF.

2.3 Topic Maps

Topic Maps [2] have 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. Topic Maps can be used as a format for the representation of multi-dimensional subject-based indices for document collections. Topic Maps can also be used as a format for interoperable knowledge representation.

The original ISO standard specified an SGML syntax for the exchange of Topic Maps. To make Topic Maps applicable on the Web, the XML Topic Maps standard has been drafted [3]. XTM defines an XML syntax for Topic Maps and gives a specific, albeit slightly simplified, data model of a Topic Map. Both the SGML syntax and the XML syntax incorporate syntax shortcuts for complex data model constructs. The XML syntax presented in [3] has been appended to [2] after publication.

Several representations on the object layer have been proposed for Topic Maps. We will adhere to the graph representation described in [6]. This graph representation knows four different kinds of arcs and three different kinds of arcs. The nodes do not differ in their properties, but in which arcs they can be connected to. Additionally, each node can have several *subject identity points*. Subject identity points serve partly serve as unique identifiers for the nodes.

The semantics of the graph nodes defined on the object layer is that of subjects, defined as anything that can be referred to in human discourse. These subjects are divided into *topics*, *associations* and *scopes*, corresponding to the three different node types. Topics can have a number of characteristics, which can be bound to them by means of associations. The processing models described in [7] and [6] state some semantic constraints on the graph, which have to be enforced in order to produce a *consistent* Topic Map. Basically, these constraints ensure that no duplicate topics occur in a consistent Topic Map.

3 Integration Approach

Our general approach is that we model a graph representation of a Topic Map with the means that an RDF graph gives us. This is an approach that has been termed "modeling the model" in [13]. In this approach, all information from the source model is preserved and just represented in another format. Thus, this transformation can also be seen as a syntax transformation. We picked this approach because it has an advantage over an approach that would perform a semantic mapping between representations. A semantic mapping will most likely incur loss of information and thus make an inverse mapping impossible. The semantic mapping approach is called “mapping the model”in [13].

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 [6] 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 [11], [15] as a means for data integration [10] [14] and transformation [4]. 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 [6]. 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 queriable.

We will now describe the different aspects of the representation of Topic Maps as RDF with respect to the layered data model described in [12].

3.2 Object layer

The representation of Topic Maps as RDF is a graph transformation on the object layer of the layered data model. The object layer describes how object identity is established and how binary relationships are described in a certain data model.

The RDF Model and Syntax description gives a graph model for RDF. The Topic Maps standard does not enforce a certain internal representation for a Topic Map. Instead, several processing models have been proposed, which describe how to deserialize an abbreviated syntax into a consistent graph-based internal data structure [7], [6]. We use the graph model presented in [6]. This graph model has the characteristics described in Section 2. Our goal is to map the Topic Map graph representation onto an RDF graph representation without any loss of information. We do this by mapping each element of the Topic Map graph described in [6] to a corresponding construct in RDF.

A prerequisite for the mapping is, that the Topic Map graph is consistent, i.e. there are no redundant elements in the Topic Map graph [6]. A Topic Map graph $tm = (N, A, S)$ consists of a set of nodes N, which have the three types a, t and s, a set of arcs A, which have the different types $associationMember$, $associationScope$, $associationTemplate$ and $scopeComponent$ and a set of resources S which indicate or constitute a subject. The *associationMember* arc has a t-node attached as a role label. Each node has at most one subject constituting resource and any number of subject indicating resources attached to it. The connection with these resources is not part of the graph [6], but taken care of in the implementation domain. However, for a mapping without loss of information, we need to consider those resources as well.

An RDF Model graph $r = (R, L, ST)$ consists of a set of resources R, a set of literals L and a set of statements ST. Our mapping m maps the set of all consistent Topic Maps TM to the set of all RDF Models R. The set of Nodes N is mapped to the set of resources R in RDF. The set of arcs A is mapped to the set of statements ST in RDF. The set of subject indicating/constituting resources S is also mapped to the set of resources R in RDF.We map the Topic Map graph to an RDF graph by first mapping the graph nodes and then mapping the arcs.

Each node in the Topic Map graph is mapped to a resource in the RDF model. The ID of the RDF resource is the ID of one of the subject identity points of the Topic Map node. If there is no subject identity point for the node, an ID is generated. For the rest of the subject identity points, statements are generated, which connect the subject identity points to its node in the RDF graph. An RDF statement is generated, which identifies the type of node that has been mapped. The Topic Map graph model knows three different kinds of nodes. We make use of the namespace capability of RDF to define the three types of nodes available in Topic Maps. The node types are defined as shown in Figure 4. An exemplary mapping of a Topic Map node to an RDF graph is shown in figure 2.

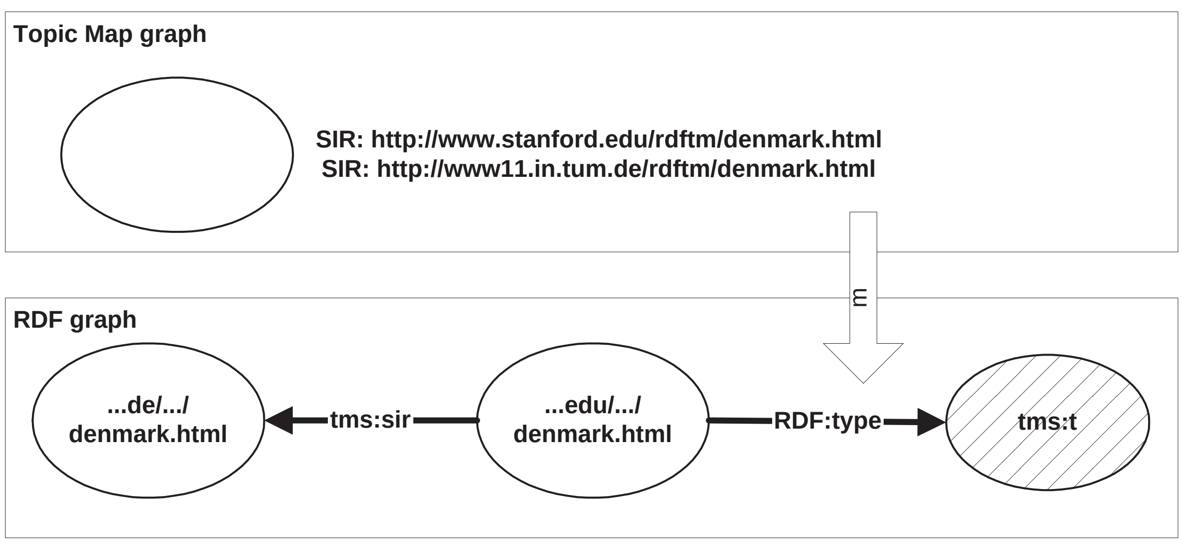

Figure 2: Exemplary mapping of a Topic Map node to an RDF graph

After all the nodes have been mapped, we map the arcs in the Topic Map graph to statements in the RDF graph. For each arc between two nodes n_1 and n_2 we generate an RDF statement. The property of the statement corresponds to the arc type in the Topic Map graph. The corresponding properties are defined in the schema in Figure 4. Although arcs in the Topic Map graph are not explicitly directed, they have an implicit directionality given through the node types at each arc end. Thus, the RDF graph is not more constrained than the Topic Map graph in that respect. If the mapped arc is an $associationMember$ arc, it has a role label in the Topic Map graph. To represent this in the RDF graph, we reify the RDF statement signifying this arc and bind the role label node to this statement with the $roleLabel$

property defined in Figure 4. The mapping of an $associationMember$ arc between two nodes is shown in Figure 3.

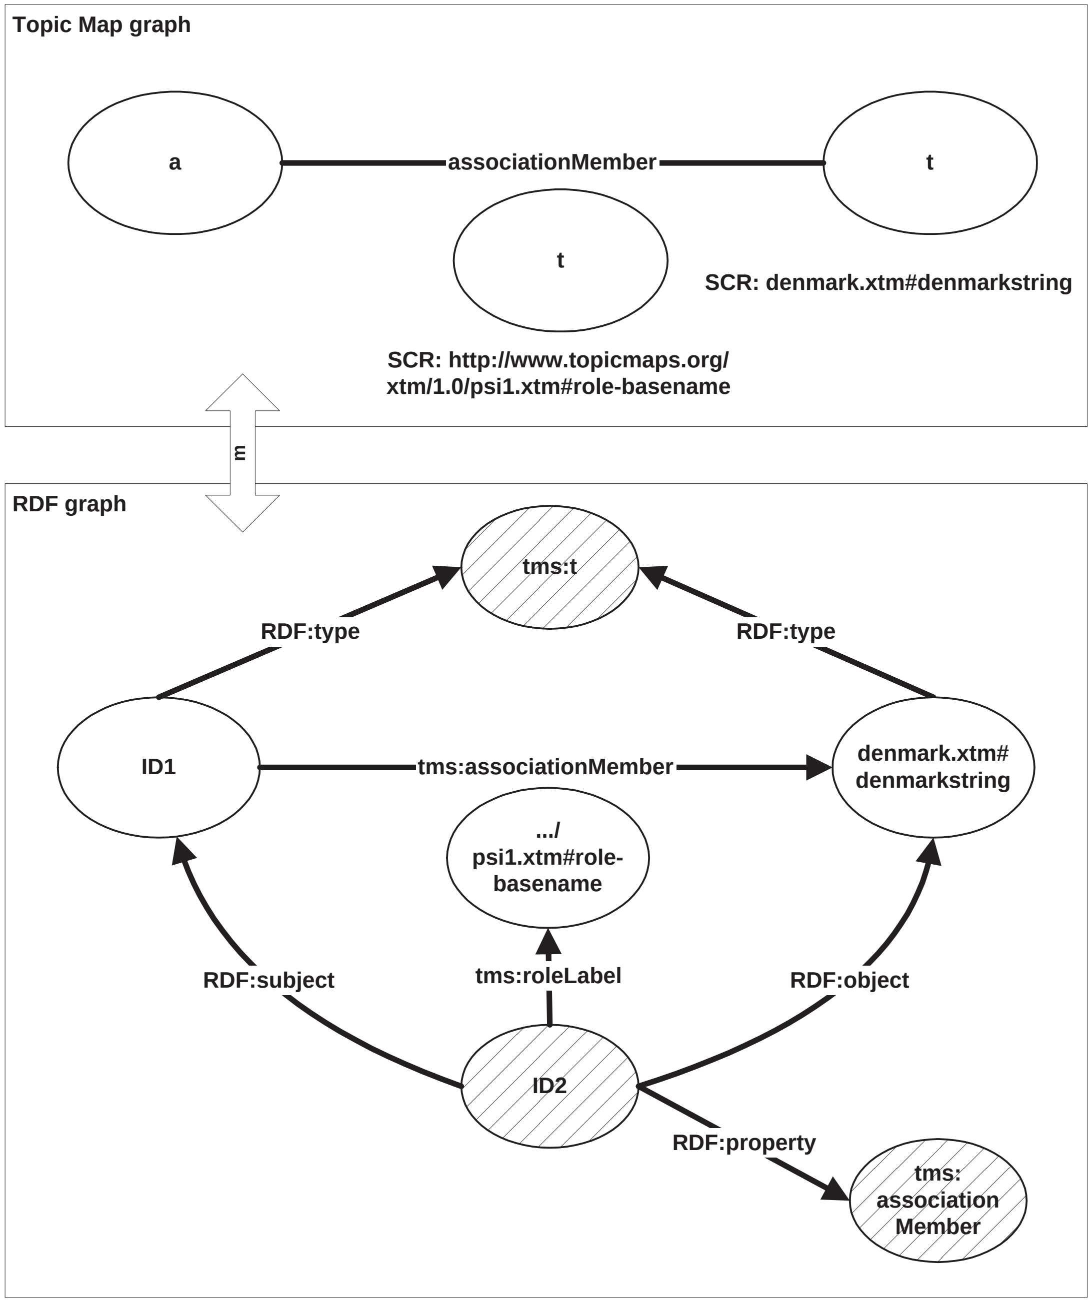

Figure 3: Exemplary mapping of a Topic Map associationMember arc to an RDF graph

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.

4 Application Example

We will now present an example for the integration of two heterogeneous information sources for querying. The first source is a Topic Map serialized in XTM [3] based on the CIA World Fact Book. The second source is the Open Directory[2], which is represented in RDF. We would like to find travel information for countries which have petroleum as a natural resource. Countries with petroleum as a natural resource can be found in the CIA World Fact Book and travel information can be found in the Open Directory collection of web pages. First, we will present how a part of the Topic Map source is mapped to RDF. Then, we will show how the two information source can be jointly queried.

4.1 Mapping the Topic Map source to RDF

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 4 shows the RDF schema definition.

```
<rdf:RDF xmlns:rdf="http://www.w3c.org/1999/02/22-rdf-syntax-ns#"
  xmlns:rdf="http://www.w3c.org/2000/01/rdf-schema#"
  xmlns:tms="http://www-db.stanford.edu/rdftmmapping/tm-schema#"
  xmlns="http://www-db.stanford.edu/rdftmmapping/tm-schema#">
  <rdfs:Class ID="t"/>
  <rdfs:Class ID="a"/>
  <rdfs:Class ID="s"/>
  <rdf:Property ID="associationMember"/>
  <rdf:Property ID="associationScope"/>
  <rdf:Property ID="associationTemplate"/>
  <rdf:Property ID="scopeComponent"/>
  <rdf:Property ID="roleLabel">
  <rdf:Property ID=="sir">
</rdf:RDF>
```

Figure 4: The 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 [6], which is required to generate a valid Topic Map graph from an abbreviated syntax. Figure 5 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 4.1.

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

[2] http://www.dmoz.org/

```
<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 5: XTM document subpart

Figure 4.1 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 base name "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 7 shows the mapped RDF graph.

It can be seen in Figure 7 that the graph can be translated in a straightforward manner. The RDF graph has additional type edges to signify the node types. All nodes in the graph which have no type edges are assumed to be of type topic in this graph. As IDs of each of the nodes we used the ID of either the respective XTM element, or generated an ID. The additional role topics, which are attached to the association member edges in the Topic Map graph, are modeled by reification of a statement in RDF: The statement that signifies the association member edge from a topic to an association is reified and becomes the subject in another statement that has the role topic as an object and the RDF-Schema-defined roleLabel as its property.

Although the mapping transforms undirected arcs into directed arcs, the mapping between the two graph representations is still a bijective mapping. The direction of arcs in the Topic Maps graph model is implicit. For querying purposes, arcs in the RDF graph of a Topic Map have to be queried in two directions.

By translating all graph constructs mentioned in the XTM processing model to an RDF

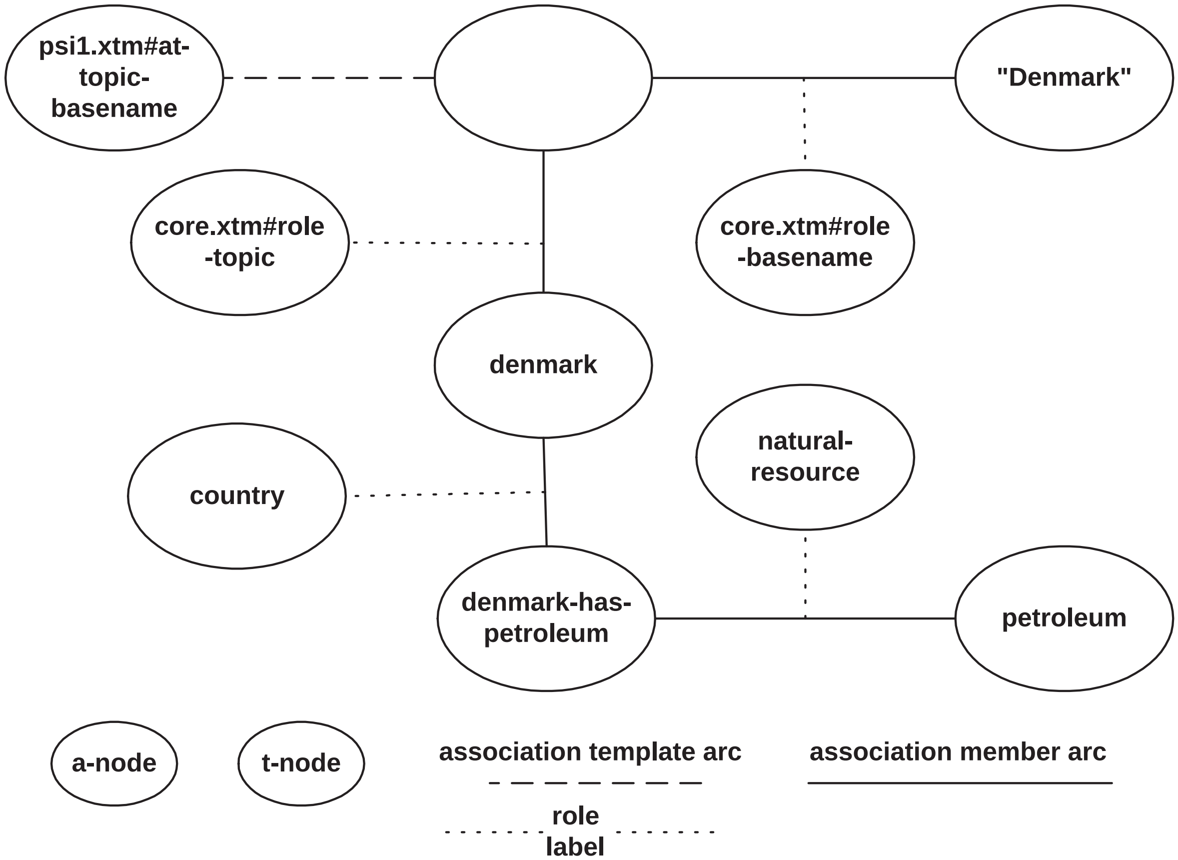

Figure 6: The generated Topic Map graph

graph we essentially generated an RDF representation of a Topic Map. We can now query this RDF graph with an RDF query language. An example for the utility of this will now be shown.

4.2 Joint querying of the information sources

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 [3], 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.

For our query example, we will assume the existence of a query engine which can query distributed information resources that are represented in RDF. The basis for such a query engine can be the query infrastructure and F-Logic syntax presented in [9]. The extension that

[3]http://www.cia.gov/cia/publications/factbook/

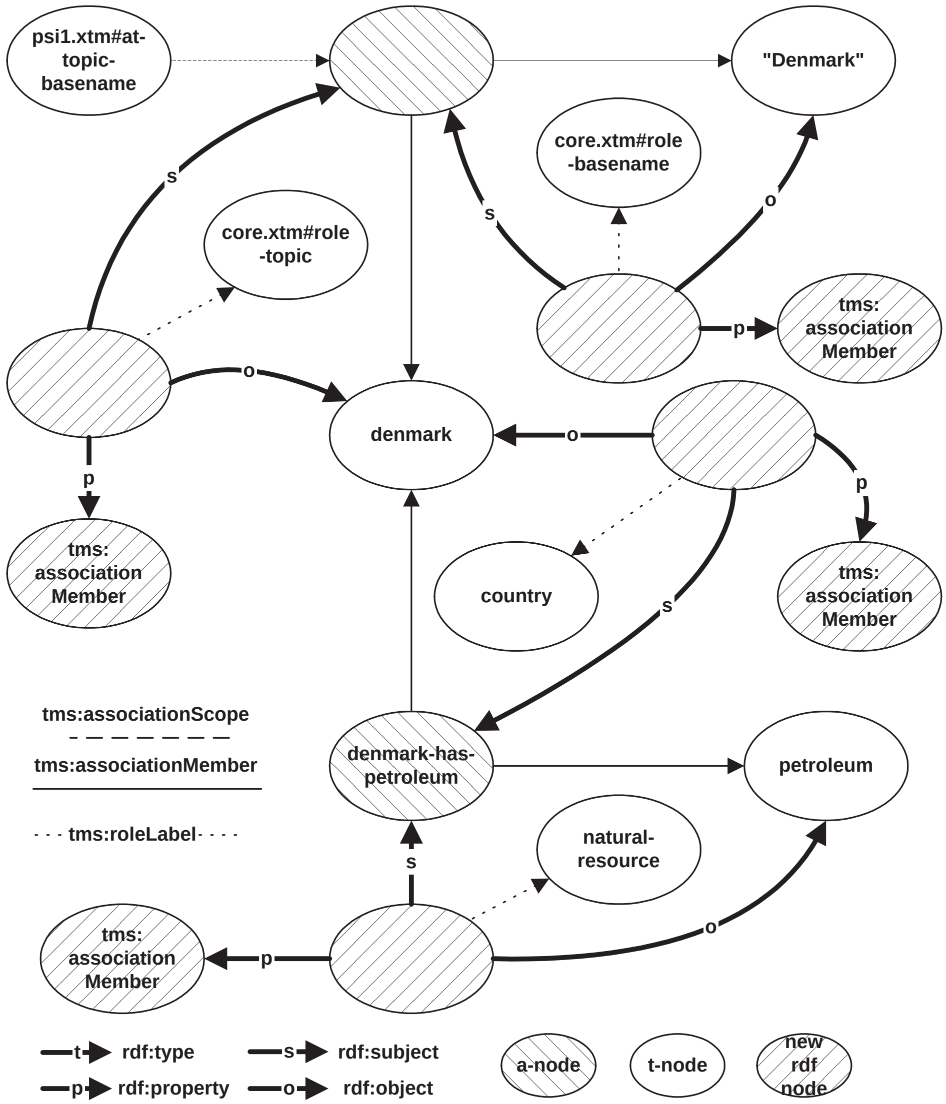

Figure 7: The generated RDF Topic Map graph

would have to be performed is to alllow to specify information sources to which certain parts of the query have to be directed. Figure 8 shows an example of a query in F-Logic syntax, as introduced in [9]. The query uses the assumed capability of the query engine to specify source with the @ character. The query in Figure 8 can now be posed to query the two above mentioned information sources.

This example query assumes the existence of a name mapping, which resolves the naming differences between resources (`mapsTo` property). Please note also that the query in Figure 8 is simplified in that naming conventions of DMOZ are not considered here. Also, the `Travel_and_Tourism` property will have to be constructed from the `DMOZCountry` URI.

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

```
FORALL pages <- Country, DMOZCountry, X, Y, 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 8: Query in F-Logic Syntax over DMOZ and RDF-based Topic Map

World Factbook Topic Map, in the RDF representation given above. Now we are able to query the DMOZ data for travel information on this country. The result of the query is a list of web pages from DMOZ categories like *Top/ Regional/ Europe/ Denmark/ Travel*, etc.

It can be seen that by representing a Topic Map in RDF, the information source becomes queriable with an RDF query language. But the actual query also requires a query processor, which can handle the distributed sources.

5 Implementation

The implementation of our RDF adapter for Topic Maps can handle the XTM syntax of Topic Maps. Both [2] and [3] constrain their normative part of the standard on the specification of an exchange syntax for Topic Maps. In order to represent a Topic Map with RDF, a graph model has to be constructed from a Topic Map document. Our implementation considers the XTM syntax and constructs a graph representation according to the processing model presented in [6]. The construction of the graph model is performed through a graph-based API proposed in [5]. 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 SAX-based parser, which feeds events to our implementation of the processing model, which then constructs the RDF graph. After constructing the graph, the redundancy rules are enforced.

6 Related Work

In [8], a general approach to integration of heterogeneous model-based information has been presented. It is shown that in principle all model based information can be represented by an RDF based meta-model. It is shown that this also includes Topic Maps. However, the authors do not go into details about this specific mapping.

In [13] two general approaches to the integration have been proposed. The first approach shows how Topic Maps can be modeled 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. The examples show the general approach of mapping Topic Maps to RDF.

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 Topic Maps. Pure syntax transformations have been proposed [4], 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.

7 Conclusion

Interoperability is of greatest importance for the future 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. Our work builds on existing work on general approaches for the integration of model based information resources. In contrast to those general approaches we showed a detailed mapping specifically from XTM Topic Maps to RDF. We achieved this by adopting an internal graph representation for Topic Maps, which has been published as part of one of the processing models for Topic Maps. 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.

References

[1] Resource Description Framework (RDF) Model and Syntax Specification, Feb. 1999. W3C Recommendation.

[2] ISO/IEC 13250: Topic Maps, Dec. 1999. ISO/IEC FCD, April 1999.

[3] XML Topic Maps (XTM) 1.0, Mar. 2001. topicmaps.org Specification, http://www.topicmaps.org/xtm/1.0/.

[4] http://lists.w3.org/Archives/Public/www-rdf-interest/2001Mar/0062.html

[4] Serge Abiteboul, Sophie Cluet, and Tova Milo. Correspondence and Translation for Heterogeneous Data. In *6th International Conference on Database Theory*, 1997.

[5] Khalil Ahmed. Developing a Topic Map Programming Model. In *Knowledge Technologies 2001*, Mar. 2001.

[6] Michel Biezunski and Steven R. Newcomb. topicmaps.net's Processing Model for XTM 1.0, version 1.0.1, May 2001. http://www.topicmaps.net/pmtm4.htm.

[7] Michel Biezunski and Steven R. Newcomb. Xml Topic Maps Processing Model 1.0, Mar 2001. http://www.topicmaps.org/xtm/1.0/xtmp1.html.

[8] Shawn Bowers and Lois Delcambre. Representing and Transforming Model-Based Information. In *International Workshop on the Semantic Web (SemWeb), in conjunction with ECDL 2000*, Sep. 2000.

[9] Stefan Decker, Dan Brickley, Janne Saarela, and Jürgen Angele. A Query and Inference Service for RDF. In *QL '98 - Query Languages Workshop*, Dec. 1998.

[10] Hector Garcia-Molina, Jan Hammer, K. Ireland, Y. Papakonstantinou, Jeff Ullman, and Jennifer Widom. Integrating and Accessing Heterogeneous Information Sources in TSIMMIS. In *AAAI Symposium on Information Gathering*, pages 61–64, Mar. 1995.

[11] Jan Hammer, J. Mc Hugh, and Hector Garcia-Molina. Semistructured Data. In *First East-European Workshop on Advances in Databases and Information Systems - ADBIS '97*, Sep. 1997.

[12] Sergey Melnik and Stefan Decker. A Layered Approach to Information Modeling and Interoperability on the Web. In *ECDL '00 Workshop on the Semantic Web*, Sep. 2000.

[13] Graham D. Moore. RDF and Topic Maps - An Exercise in Convergence. In *XML Europe 2001*, May 2001.

[14] Y. Papakonstantinou, Hector Garcia-Molina, and Jennifer Widom. Object Exchange Across Heterogeneous Information Sources. In *ICDE '95*, Mar. 1995.

[15] Dan Suciu. An Overview of Semi-Structured Data. *SIGACTN: SIGACT News*, 29:pp. 28–38, 1998.