Metamorphosis – An Environment to Achieve Semantic Interoperability with Topic Maps

Giovani Rubert Librelotto¹ and José Carlos Ramalho² and Pedro Rangel Henriques²

¹UNIFRA – Centro Universitário Franciscano Rua dos Andradas, 1614 – 97010-032 – Santa Maria – RS – Brazil

²Universidade do Minho, Departamento de Informática Braga, Portugal, 4710-057

giovani@unifra.br, {jcr,prh}@di.uminho.pt

Abstract. Nowadays, data handled by an institution or company is spread out by more than one database and lots of documents of different types. To extract the information implicit in that data, it is necessary to pick parts from those various archives. To obtain a general overview, those information slices should be gather. Different approaches can be followed to achieve that integration, ranging from the merge of resources till the fusion of the extracted parts. In this paper, we introduce Metamorphosis – a Topic Maps oriented environment to generate conceptual navigators for heterogenous information systems – and we argue that Metamorphosis can be used to achieve the referred interoperability.

1. Introduction

Daily, a lot of data is produced by every institution or company. To satisfy the storage requirements, these organizations use most of the times relational databases, which are quite efficient to save and to manipulate structured data. Unstructured data (appearing inside documents) is stored in plain or annotated text files.

There is a problem when these organizations require an integrated view of their heterogeneous information systems. It is necessary to query/exploit every data source, but the access to each information system is different. In this situation, there is a need for an approach that extracts the information from those resources and fuses it. Usually this is achieved either by extracting data and loading it into a central repository that does the integration before analysis, or by merging the information extracted separately from each resource into a central knowledge base.

Topic Maps [Park and Hunting 2003] are a good solution to organize concepts, and the relationships between those concepts, because they follow a standard notation – ISO/IEC 13250 [Biezunsky et al. 1999] – for interchangeable knowledge representation. We are using successfully, for some years, this technology for classification and integration of documents in the area of digital archiving.

However, the process of ontology development based on topic maps is complex, time consuming, and it requires a lot of human and financial resources, because they can have a lot of topics and associations, and the number of resources can be very large.

To overcome this problem, we developed *Metamorphosis*. *Metamorphosis* makes possible the Topic Maps extraction, validation, storage, and browsing. It is composed of

three main modules: (1) *Oveia* extracts data, from heterogeneous information systems, according to an ontology specification, and stores it in a topic map; (2) *XTche* validates the generated topic map, according to a constraint specification; (3) *Ulisses* browses the topic map, enabling a conceptual navigation and query over the resources.

The remainder of the paper is structured in the following sections: in section (sec.2) will introduce *Metamorphosis*, then a description of each module is presented with some detail (*Oveia* in sec.3, *XTche* in sec.4 and *Ulisses* in sec.5). Before concluding remarks (sec.7) we compare our proposal with related work (sec.6).

2. Metamorphosis: an environment to deal with Topic Maps

Topic Maps are very well suited to represent ontologies [Wrightson 2001]. Ontologies play a key role in many real-world knowledge representation applications, and namely the development of Semantic Web.

An ontology is a way of describing a shared common understanding, about the kind of objects and relationships which are being talked about, so that communication can happen between people and application systems [Guarino 1998]. In other words, it is the terminology of a domain (it defines the universe of discourse). As a real example consider the thesaurus used to search in a set of similar, but independent, websites.

The ability of Topic Maps to link resources and to organize these resources according to a single ontology, will make Topic Maps a key component of the new generation of Web-aware knowledge management solutions. In addition, the growing repertoire of techniques for simplifying, merging and interrelating ontologies can be used to combine or articulate Topic Maps representing different ontologies, thus enabling different sets of information resources to be used together in a controlled and scalable way [Freese 2000].

The main idea behind *Metamorphosis* is to integrate the specification of conceptual networks or ontologies, with their storage and navigation, as well as, their automatic creation and validation.

One of the first *Metamorphosis*' applications was the production of website maps for conceptual navigation; another of our former concerns was the contents publishing in the context of e-learning. *Metamorphosis* can be also used to test some functionalities of a dynamic web system because it creates, in a fast way, a web interface that interacts directly with data sources.

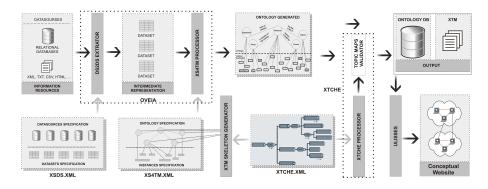


Figure 1. Metamorphosis Architecture

Figure 1 shows *Metamorphosis*' architecture that came up from the principles underlying our proposal. This architecture is composed of:

- (1) **Information Resources:** It is composed of the data sources: XML documents, databases, Web pages, etc.
- (2) **XSDS and XS4TM specifications:** They are domain specific languages to define the topic map extraction.
- (3) Oveia: The processor that builds topic maps. Its core is a processor that extracts the topics instances from the information resources and builds a topic map. It reads and processes the XSDS and XS4TM specifications.
- **(4) Generated topic map:** The topic map automatically generated by *Oveia* stored as an XTM file or alternatively a relational database.
- (5) **XTche specification:** A topic maps constraint specification language based on TMCL (Topic Maps Constraint Language) [Nishikawa and Moore 2003] that allows to define rules for topic maps semantic validation.
- **(6) XTche Processor:** The processor that consumes the previous XTM file and verifies the topic map according to a set of constraints defined in *XTche* language.
- (7) **Valid topic map:** The previous topic map automatically validated by *XTche*.
- (8) Ulisses: The processor that takes a topic map and produces a whole semantic Web site, a set of Web pages where it is possible to navigate through structural or syntactic links as well as through a network of concepts.
- (9) Conceptual Web site: It is the generated Web site that allows the semantic navigation over the topic map extracted from information resources.

In the next sections we are going to discuss the main pieces of this architecture: *Oveia*, *XTche*, and *Ulisses*, in order to demonstrate how the overall system can accomplish the task we have stated at the beginning.

3. Oveia

The ontology extractor – *Oveia* – is based on ISO/IEC 13250 Topic Maps [Biezunsky et al. 1999]. *Oveia* extracts information fragments from heterogeneous information systems according to an XSDS specification and builds the topic map according to an ontology specified in XS4TM language.

The *Oveia* architecture is shown in Figure 2 and it is composed mainly of five components. The dataset extractor receives an XSDS specification – providing metadata about the physical data sources that will be used to query each source in order to get the data needed for the ontology construction – and generates the intermediate representation (called *datasets*) – containing the data extracted from resources. The XS4TM processor takes as input these datasets and an XS4TM specification generating a topic map, in an internal format. An output generator stores the topic map in an OntologyDB or in an XTM file. The following subsections describe this architecture in detail.

3.1. XSDS — XML Specification for Data Sources

Oveia supports the concept of extraction drivers. A driver extracts data from a data source and stores it in an intermediate representation, called *datasets*. XSDS language defines the transformations and filters over the data sources. XSDS gives precise information about each data source that should be scanned to extract topics and associations.

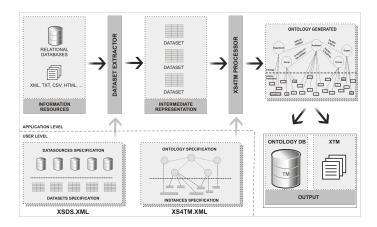


Figure 2. Oveia Architecture

An XSDS specification has two parts: *datasources* and *datasets*. The first one defines the path to the physical resources. This part has a set of attributes that indicates which extraction driver will be used and provides values for the corresponding parameters. The last one declares which data (record fields or DTD elements) must be extracted from each *datasource*. A *datasource* can be used to specify the extraction of several *datasets*.

3.2. Datasets: Intermediate Representation

The *datasets* compose the intermediate representation that contains the extracted data from the resources. Each *dataset* has a relation to an entity in these resources and it is represented through a table, where each line is a record following the structure specified in XSDS. The *datasets* representation guarantees that *Oveia* sees an uniform data structure that represents all the participating resources.

The *dataset* declaration is composed by a query to extract the data from the resources. Each dataset has an unique identifier that will be used throughout the architecture to reference a particular *dataset*. The *datasets* are very simple, while providing the expressive power and flexibility needed for integrating information from disparate sources.

The *Dataset Extractor*¹ is composed of several extraction drivers (at moment, two), each one responsible for handling a specific type of source. The driver uses the appropriate technology to make the connection (e.g. JDBC – Java DataBase Connectivity – for databases, and an XML parser for annotated documents), and then the extraction of data is expressed in the query language adequate to the type of source in use: SQL will be used to extract information from a relational database while XPath will be used for the extraction in XML documents. Finally, the extracted data is stored in the *datasets*.

3.3. XS4TM — XML Specification for Topic Maps

XS4TM is a domain specific language conceived to specify the process of ontology extraction from information systems; in our case, from the *datasets* representation.

Looking at a topic map an ontology designer can think of it as having two distinct parts: an ontology and an object catalog (instances). The ontology is defined by topic

¹A processor that scans the input data sources to get desired data into the *datasets*, in agreement with an XSDS specification.

types, association types, occurrence types, role types, etc. The catalog is composed of a set of pointers to information objects that are present in the resources and are linked to the ontology. So, a specification in XS4TM is composed of two parts:

Ontology: the definition of the ontology requires in XS4TM the same effort as in XTM; it is necessary to specify every topic type, association type, occurrence type, ...; **Instances:** the instances definition describes each topic and association that will be extracted from the intermediate representation.

The XS4TM Context Free Grammar is based in XTM 1.0 [Pepper and Moore 2001]. The *ontology* and *instances* elements have the same syntax as the *topicMap* element in XTM model.

3.4. XS4TM Processor

This component uses the XS4TM specification and retrieves the information it needs to build the ontology from the *datasets*. It is an interpreter that takes advantage of the information organization in datasets (an internal universal representation for extracted data) and generates all the associations between the relevant topics according to XS4TM.

The XS4TM processor's behavior can be described in three steps: reads the the XS4TM specification and extracts from the datasets the topics and associations found; creates the topic map; finally, stores it into an database or an XTM file.

4. XTche – A Topic Maps Constraint Language

When developing real topic maps, it is highly convenient to use a system to validate it; this is, to verify the correctness of the actual instance against the formal specification of that family of topic maps (according to the intention of its creator).

So, a specification language that allows us to define the schema and constraints of a family of Topic Maps is necessary. A list of requirements for the new language was recently established by the ISO Working Group – the ISO JTC1 SC34 Project for a Topic Map Constraint Language (TMCL) [Nishikawa and Moore 2003]. *XTche* language meets all the requirements in that list.

XTche [Librelotto et al. 2005] is designed to allow users to constrain any aspect of the topic map; for instance: topic names and scopes; association members, roles and players allowed in an association, instances of a topic (enumeration), association in which topics must participate, occurrences cardinality, etc.

Like XTM, XTche specifications can be too verbose; that way it is necessary to define constraints in a graphical way with the support of a visual tool. To overcome this problem, XTche syntax follows the XML Schema syntax; so, any XTche constraint specification can be written in a diagrammatic style with a common XML Schema editor. At the end the textual output of that edition (XML Schema code) should be processed to obtain a TMValidator.

4.1. XTche Processor and TM-Validator

A XTche specification, listing all the conditions (involving topics and associations) that must be checked, specifies the Topic Map validation process (TM-Validator), enabling

the systematic codification (in XSL) of this verification task. In that circumstances we understood that it was possible to generate automatically the validator developing another XSL processor to translate an *XTche* specification into the TM-Validator XSL code.

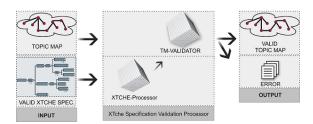


Figure 3. XTche Validation Process

According to Figure 3, the XTche processor is the TM-Validator generator; it takes a topic map constraint specification (an XML-Schema, written according to the *XTche* language), and generates an XSL stylesheet (the TM-Validator) that will process an input topic map in order to verify its correctness.

5. Ulisses

Ulisses can be seen as a website generator from a XTM document (the "source" topic map) – this explains why we decided to integrate it as the last layer of *Metamorphosis*. It was conceived to be a autonomous (it can be used outside of *Metamorphosis* context) and simple way of creating full sites, with design, content and topical links; however, the layout of the site generated can be customized (page design, colors, ...) to satisfy the specific user needs. Allowing the navigation on a conceptual network (an ontology described by the source topic map), *Ulisses* can be seen as a useful tool to develop the so called *semantic web*.

The basic idea behind the website generation is to create one *HTML page* for each topic or association. Hyperlinks are then used to connect related topics or topics and associations. A navigation menu, allowing to go back to the home page or to choose another view of the topic map, is always present in every page.

As told above, each topic or association name displayed in one *HTML page* is a hyperlink to the respective page, thus implementing the conceptual navigation over the semantic network described by the topic map.

We developed three different versions of *Ulisses: Ulisses I and II* read the input from a XTM text file, while *Ulisses III* takes as input a *OntologyDB* (see above, sec. 3). Concerning the generation strategy, the original version (*Ulisses I*) is a *static generator*—it processes just once the XTM file and creates at that time all the website pages; the generation is time-consuming and the site directory huge, however the topic map navigation is very fast. The drawback of that approach is that any change on the "source" TM implies the complete regeneration; otherwise the navigator becomes inconsistent/obsolete. To overcome that problem, the other two versions follow the opposite approach, implementing a *dynamic generation*; the first page (the homepage) is created at generation time and the others are created by need at navigation time.

6. Related Work

In terms of related work [Wache et al. 2001] we did not found an environment that can be compared to *Metamorphosis*. So, the comparisons below are among the main *Metamorphosis*' modules and their related work.

TSIMMIS[Rys 1998] is a project that aims to provide tools for accessing, in an integrated fashion, multiple information sources, and to ensure that the information obtained is consistent. TSIMMIS gives a centralized view of the information that is distributed in the information system. *Oveia* was developed to allow a conceptual navigation over the heterogeneous information systems. This conceptual navigation is driven by an ontology specified from metadata extracted from information systems.

In another comparison, KAON REVERSE [Volz 2003] has advantages concerning the use of a graphical interface for the specification of the ontology against *Oveia*. It also allows the use of reverse engineering of data sources to help creating the mapping. On the other side, *Oveia* is more flexible concerning data source formats and the specification process. To represent the ontology, KAON REVERSE adopts RDF; *Oveia* generates ontologies and stores them in an ontology database (*OntologyDB*) or in an XTM file.

AsTMa! [Barta 2003] is another Topic Maps constraint language that has a mechanism to validate a topic map document against a given set of rules, like *XTche* language. That language has logic operators like NOT, AND and OR, simple logical quantifiers, and regular expressions.

When a comparison between *XTche* and the related works is done, some advantages is detected: *XTche* has a XML Schema-based language, a well-known format. In addition, *XTche* allows the use of an XML Schema graphical editor, like XMLSpy. With the diagrammatic view, it is easy to check visually the correctness of the specification.

7. Conclusion

This paper describes the integration of heterogeneous information systems using the ontology paradigm, in order to generate an homogeneous view of these resources.

Metamorphosis let us achieve the semantic interoperability among heterogeneous information systems because the relevant data, according to the desired information specified through an ontology, is extracted and stored in a topic map. The environment validates it against a set of rules defined in a constraint language. That topic map provides information fragments (the data itself) linked by specific relations to concepts at different levels of abstraction. Note that not all data items need to be extracted from the sources to the Topic Map. We only extract the necessary metadata to build the intended ontology. This ontology will have links to enable a browser to access all data items.

Thus the navigation over the topic map is led by a semantic network and provides an homogeneous view over the resources – this justifies our decision of call it semantic interoperability [Mitra and Wiederhold 2001].

Although developed for use in our main working area – XML documents processing applied to Public Archives and Virtual Museums – we are convinced that *Metamorphosis* can be applied with similar success in the general area of information system for data integration, analysis, and knowledge exploitation.

References

- Barta, R. (2003). AsTMa! Bond University, TR. http://astma.it.bond.edu.au/constraining.xsp.
- Biezunsky, M., Bryan, M., and Newcomb, S. (1999). ISO/IEC 13250 Topic Maps. ISO/IEC JTC 1/SC34. http://www.y12.doe.gov/sgml/sc34/document/0129.pdf.
- Freese, E. (2000). Using Topic Maps for the representation, management and discovery of knowledge. In *XML Europe 2000 Proceedings*. http://www.gca.org/papers/xmleurope2000/papers/s22-01.html.
- Guarino, N. (1998). Formal Ontology and Information Systems. In *Conference on Formal Ontology (FOIS98)*. http://www.ladseb.pd.cnr.it/infor/Ontology/Papers/FOIS98.pdf.
- Librelotto, S. R., Ramalho, J. C., and Henriques, P. R. (2005). Constraining topic maps: a TMCL declarative implementation. In *Extreme Markup Languages 2005:**Proceedings. IDEAlliance. http://www.mulberrytech.com/Extreme/
 Proceedings/html/2005/Ramalho01/EML2005Ramalho01.html.
- Mitra, P. and Wiederhold, G. (2001). An algebra for semantic interoperability of information sources. In *BIBE '01: Proceedings of the 2nd IEEE International Symposium on Bioinformatics and Bioengineering*, page 174, Washington, DC, USA. IEEE Computer Society.
- Nishikawa, M. and Moore, G. (2003). Requirements for a Topic Map Constraint Language JTC 1 NP Number. ISO/IEC 19756. ISO/IEC JTC 1/SC34 N0405. http://www.y12.doe.gov/sqml/sc34/document/0405.htm.
- Park, J. and Hunting, S. (2003). *XML Topic Maps: Creating and Using Topic Maps for the Web*, volume ISBN 0-201-74960-2. Addison Wesley.
- Pepper, S. and Moore, G. (2001). XML Topic Maps (XTM) 1.0. TopicMaps.Org Specification. http://www.topicmaps.org/xtm/1.0/.
- Rys, M. (1998). TSIMMIS The Stanford-IBM Manager of Multiple Information Sources. http://www-db.stanford.edu/tsimmis/tsimmis.html.
- Volz, R. (2003). KAON REVERSE. http://kaon.semanticweb.org/ alphaworld/reverse/view.
- Wache, H., Vögele, T., Visser, U., Stuckenschmidt, H., Schuster, G., Neumann, H., and Hübner, S. (2001). Ontology-based integration of information a survey of existing approaches. In Stuckenschmidt, H., editor, *IJCAI–01 Workshop: Ontologies and Information Sharing*, pages 108–117.
- Wrightson, A. (2001). Topic Maps and Knowledge Representation. Ontopia. http://www.ontopia.net/topicmaps/materials/kr-tm.html.