



I4-D1

Survey over Existing Query and Transformation Languages

Project number:	IST-2004-506779
Project title:	Reasoning on the Web with Rules and Semantics
Project acronym:	REWERSE
Document type:	D (deliverable)
Nature of document	R (report)
Dissemination level:	PU (public)
Document number:	IST506779/Munich/I4-D1/D/PU/a1
Responsible editor(s):	Tim Furche
Reviewer(s):	Wolfgang May
Contributing participants:	Hannover, Heraklion, Manchester, Munich, Nancy, Venice, webXcerpt
Contributing workpackages:	I4
Contractual date of delivery:	31 August 2004

Abstract

A widely acknowledged obstacle for realizing the vision of the Semantic Web is the inability of many current Semantic Web approaches to cope with data available in such diverging representation formalisms as XML, RDF, or Topic Maps. A common query language is the first step to allow transparent access to data in any of these formats. To further the understanding of the requirements and approaches proposed for query languages in the conventional as well as the Semantic Web, this report surveys a large number of query languages for accessing XML, RDF, or Topic Maps. This is the first systematic survey to consider query languages from all these areas. From the detailed survey of these query languages, a common classification scheme is derived that is useful for understanding and differentiating languages within and among all three areas.

Keyword List

reasoning, query language, XML, RDF, Topic Maps, OWL, classification, Semantic Web

Survey over Existing Query and Transformation Languages

Tim Furche¹, François Bry², Sebastian Schaffert³, Renzo Orsini⁴, Ian Horrocks⁵, Michael Kraus⁶, and Oliver Bolzer⁷

¹ Institute for Informatics, Ludwig-Maximilians-Universität München
Email: Tim.Furche@ifi.lmu.de

² Institute for Informatics, Ludwig-Maximilians-Universität München
Email: Francois.Bry@ifi.lmu.de

³ Institute for Informatics, Ludwig-Maximilians-Universität München
Email: Sebastian.Schaffert@ifi.lmu.de

⁴ Dipartimento di Informatica, Università Ca' Foscari Venezia
Email: orsini@dsi.unive.it

⁵ Department of Computer Science, University of Manchester
Email: horrocks@cs.man.ac.uk

⁶ Institute for Informatics, Ludwig-Maximilians-Universität München
Email: Michael.Kraus@ifi.lmu.de

⁷ Institute for Informatics, Ludwig-Maximilians-Universität München
Email: Oliver.Bolzer@stud.ifi.lmu.de

8 September 2004

Abstract

A widely acknowledged obstacle for realizing the vision of the Semantic Web is the inability of many current Semantic Web approaches to cope with data available in such diverging representation formalisms as XML, RDF, or Topic Maps. A common query language is the first step to allow transparent access to data in any of these formats. To further the understanding of the requirements and approaches proposed for query languages in the conventional as well as the Semantic Web, this report surveys a large number of query languages for accessing XML, RDF, or Topic Maps. This is the first systematic survey to consider query languages from all these areas. From the detailed survey of these query languages, a common classification scheme is derived that is useful for understanding and differentiating languages within and among all three areas.

Keyword List

reasoning, query language, XML, RDF, Topic Maps, OWL, classification, Semantic Web

Contents

1	Introduction	1
1.1	Selection of Evaluation Criteria	2
1.2	Selection of Surveyed Query Languages	4
2	Preliminaries	7
2.1	Collection of Sample Data	9
3	Evaluation Criteria	15
3.1	Ease of Use	15
3.2	Functionality	18
3.2.1	Supported Query Types	18
3.2.2	Adequacy	24
3.2.3	Evolution and Reactivity	30
3.3	Semantics	30
3.4	Formal Properties and Implementation	31
3.5	Reasoning	32
3.6	Ontology Awareness	34
3.6.1	Type system	35
4	Query Languages for the Web: An Overview	37
4.1	XML Query Languages	37
4.2	Textual XML Query Languages	37
4.2.1	Navigational Languages	37
4.2.1.1	Lorel	37
4.2.1.2	XPath	38
4.2.1.3	XQL	39
4.2.1.4	XSLT	39
4.2.1.5	XQuery/Quilt	40
4.2.1.6	FXT	41
4.2.1.7	XPathLog	41
4.2.1.8	CXQuery	42
4.2.2	Positional Languages	42
4.2.2.1	XML-QL	42
4.2.2.2	UnQL	42
4.2.2.3	XML-RL	43
4.2.2.4	XMAS	43

4.2.2.5	XET/XDD	44
4.2.2.6	Xcerpt	44
4.3	RDF Query Languages	46
4.3.1	SquishQL-family	46
4.3.1.1	SquishQL	46
4.3.1.2	rdfDB Query Language	46
4.3.1.3	RDQL	47
4.3.1.4	BRQL	50
4.3.1.5	TriQL	51
4.3.2	Query Languages influenced by XPath, XSLT or XQuery	52
4.3.2.1	XQuery for RDF: “The Syntactic Web” Approach	52
4.3.2.2	XsRQL: An XQuery-style RDF Query Language	54
4.3.2.3	XSLT for RDF: TreeHugger and RDF Twig	54
4.3.2.4	RDFT and Nexus Query Language: XSLT-style RDF Query Languages	56
4.3.2.5	XPath-syle Access to RDF: RDF Path, RPath and RxPath	58
4.3.2.6	Versa	61
4.3.3	RQL-family	65
4.3.3.1	RQL	65
4.3.3.2	SeRQL	71
4.3.3.3	eRQL	74
4.3.4	Query Languages using a Controlled Natural Language	75
4.3.4.1	Metalog	75
4.3.5	Others	77
4.3.5.1	Algae	77
4.3.5.2	iTQL	80
4.3.5.3	N3QL	80
4.3.5.4	PerlRDF Query Language	82
4.3.5.5	R-DEVICE Deductive Language	82
4.3.5.6	RDF-QBE	83
4.3.5.7	RDFQL	84
4.3.5.8	TRIPLE	86
4.3.5.9	WQL	89
4.4	Topic Maps Query Languages	90
4.4.1	tolog: Logic Programming for Topic Maps	90
4.4.2	AsTMA?: Functional-style Querying of Topic Maps?	94
4.4.3	Toma: Querying Topic Maps inspired by SQL	95
4.4.4	Path-based Access to Topic Maps	96
4.4.4.1	XTMPPath	96
4.4.4.2	TMPPath	97
4.5	OWL Query Languages	98
4.5.1	OWL-QL	98
5	Evaluation Results	101
5.1	A Classification Scheme for Web Query Languages	101
5.2	Observations on the State of the Art of Web Query Languages	104
6	Conclusion and Outlook	107

A	A Brief History of RDF Serialization Formats	109
A.1	Introduction	109
A.2	RDF/XML: The W3C Recommendation	109
A.3	Simplified Syntaxes for RDF/XML	112
A.3.1	Unstriped Syntax	112
A.3.2	Simplified Syntax	113
A.3.3	XMP	113
A.3.4	Normalized RDF	113
A.3.5	RXML	114
A.4	Plain-Text Formats	115
A.4.1	Notation 3	115
A.4.2	N-Triples	116
A.4.3	Quads	116
A.4.4	Turtle	116
A.4.5	TriG	117
A.5	Triple-based XML Formats	117
A.5.1	RPV	118
A.5.2	TriX	119
A.5.3	RXR	120
A.6	Features Overview	121
A.7	Genealogy	121
A.8	Conclusions	122
B	Evaluation Tables	123

Chapter 1

Introduction

The “Semantic Web” is an endeavor which Tim Berners-Lee, the father of HTML and of HTTP, James Hendler, and Ora Lassila initiated in 2001 with an article in the Scientific American [37]. The “Semantic Web” vision is that of the current Web which consists of (X)HTML and documents in other XML formats being extended with meta-data specifying the meaning of these documents in forms usable by both, human beings and computers:

The Semantic Web will bring structure to the meaningful content of Web pages, creating an environment where software agents roaming from page to page can readily carry out sophisticated tasks for users. [37]

One might see the Semantic Web meta-data added to today’s Web as semantic indices similar to encyclopedias. A considerable advantage over conventional encyclopedias printed on paper is that the relationships expressed by Semantic Web meta-data can be followed by computers, very much like hyperlinks can be followed by programs, and be used for drawing conclusion using automated reasoning methods.

For the Semantic Web to function, computers must have access to structured collections of information and sets of inference rules that they can use to conduct automated reasoning. [37]

A number of formalisms have been proposed in recent years for representing Semantic Web meta-data, e.g., RDF [197], Topic Maps [144], and OWL [24]. Whereas RDF and Topic Maps provide merely a syntax for representing assertions on relationships like “a text is authored by some person”, schema or ontology languages such as RDFS [54] and DAML+OIL [140] allow to state properties of the terms used in such assertions, e.g., that no “person” can be a “text”. Building upon descriptions of resources and their schemata (as detailed in the architectural road map for the Semantic Web [34]), rules expressed in, e.g., SWRL [138] or RuleML [44], allow the specification of actions to be taken, knowledge to be derived, or constraints to be enforced.

Essential for realizing this vision is the integrated access to all kinds of data represented in any of these representation formalisms or even in standard Web languages such as (X)HTML, SVG, or any other XML format. Considering the large amount and the distributed storage of data available already on the Web, the efficient and convenient access to such data becomes *the* enabling requirement for the Semantic Web vision. It has been recognized that a reasonably

high-level, declarative query language is needed for such efficient and convenient access, as it allows to separate the actual data storage from the view of the data a query programmer operates on.

Therefore, the aim of this survey is to provide an overview over the languages considered for each of the major representation formalisms used in the nowadays Web, viz., for XML, RDF, Topic Maps, and OWL. This overview is intended to be valuable for comparing, e.g., RDF query languages among themselves, but to also provide insight on the question, whether a common query language for these representation formalisms is reasonable. Therefore, the following three questions stand at the heart of this survey

- What are the **capabilities of a query language** considered essential for the different areas? Is it possible to identify **common shortcomings** of existing approaches in an area, in particular by means of a **comparison of the issues addressed in that particular area with the other areas investigated**?
- Enabling reasoning, i.e., the ability to derive new knowledge from existing knowledge in a systematic way, is perhaps the most distinguishing feature of the “Semantic Web” vision. Convenient and effective querying in such a setting is likely to require at least some degree of reasoning abilities (e.g., for mediation of data described with differing but convertible vocabularies). Therefore this survey closely investigates, **what reasoning abilities** the query languages offer and **how these reasoning abilities are realized**.
- Indeed, the extent and realization of reasoning abilities proves to be a crucial differential for answering the question, **how to classify the query languages surveyed in this work**. In Section 5.1 a common classification scheme for Web query and transformation languages oriented on their Semantic Web “fitness” is proposed and its usefulness for understanding the differences among the query languages is demonstrated.

To this end, this survey starts in Chapter 2 with a concise introduction into the three representation formalisms considered here, viz. XML, RDF, and Topic Maps. Note, that for most of the discussion OWL is not considered separately, but rather in conjunction with RDF, since there is a number of query languages for RDF that use information represented in (some subset of) OWL for querying. However, in Section 4.5 an approach for querying ontologies represented in OWL is discussed to illustrate the challenges one faces when more powerful ontology languages are considered. Chapter 2 also introduces the scenario used in most of the query language descriptions. A small collection of data about books and their classification is introduced on an abstract level and carefully crafted representations in XML, RDF, and Topic Maps are proposed and discussed.

1.1 Selection of Evaluation Criteria: How to evaluate a Web Query Language?

Based upon this collection of sample data, Chapter 3 proposes (a) an **exhaustive set of evaluation criteria** based on requirements and use cases for Web query languages previously identified in [176, 203, 69, 115, 19, 81, 62] and extended by additional criteria for investigating the Semantic Web “fitness” of the query languages in question. As a means for better illustrating the capabilities of the query languages, the taxonomy of queries proposed in [176] is adapted

to the Semantic Web setting and (b) a **small set of queries covering each of the classes in the query taxonomy** is proposed. To compare query languages among different representation formalisms, the queries are presented on a rather abstract level, allowing them to be applied on XML, RDF, or Topic Maps data.

This combined approach has a number of merits compared to previous surveys of Web query languages, cf. [4, 101, 48, 46, 173] (surveys of XML query languages) and [56, 174, 219, 82, 131, 221] (surveys and comparisons of RDF query languages), that have mostly been based on a set of exemplary queries and limited to a small number of evaluation criteria (represented in these queries). However, the crucial aspect of the “user experience” of a language, i.e., how convenient and effective the use of a query language is for solving practical problems, is *not* fully covered by this approach, since this aspect is hard to measure without extensive experimental studies involving users with varying background and expertise. Another noteworthy limitation of this survey is that, for time and space reasons, *not all languages could be covered in the same detail*. Instead, quite a number of languages judged particularly interesting or innovative by the authors of this survey are discussed in more detail, whereas other languages presented highlighting only the most interesting features. However, for all 73 languages the full set of 111 criteria have been evaluated and gathered in tabular form in Appendix B. This limitation is rooted partially in the fact, that there has been a virtual surge of new Web query languages, in particular of Semantic Web query languages (i.e., so far mostly RDF and Topic Maps query languages) in the last two years. Since the beginning of 2004, a dozen new languages have been proposed or existing proposals have been significantly altered. This demonstrates both the high relevance and the relative immaturity of the area of (Semantic) Web query languages. Another area where this survey might be further improved in the future is on the theoretical foundations of the languages considered. Regarding, e.g., formal semantics and data model, only rather general statements are noted here. This is motivated by the lack of consideration of such issues in the majority of the language proposals forming the base of this survey.

Despite these shortcomings the approach taken in this survey also exhibits a number of advantages in contrast to previous surveys of Web query languages:

- considering both query languages for standard and Semantic Web allows a better understanding of what the crucial aspects of Semantic Web query languages might be;
- the far larger number of approaches considered gives the results a broader foundation and applicability;
- the evaluation criteria are, where possible, restricted to easily verifiable properties of the query languages;
- focusing the discussion on a selected set of languages allows more details and a better understanding for that languages.

Finally it should be noted that, as with any such survey, the selection of both criteria and sample queries is certainly subjective and might be biased towards a certain result or language, although the authors of this survey tried carefully to eliminate such bias as far as possible.

1.2 Selection of Surveyed Query Languages: What is a Web query language?

Chapter 4 presents the query languages surveyed in this paper grouped by the underlying representation formalism and, where possible, language “family” (i.e., closely related languages are discussed together to ease the understanding of commonalities and differences).

The selection of languages requires some justification. One basic premise guided this selection: Although the distinction is not always clear, the survey should focus on **languages designed primarily for providing efficient and effective access to data**. This rather narrow basic premise excludes in particular three types of languages that are also sometimes considered query languages or at least related to query languages:

- *Full programming languages and libraries or APIs for accessing XML*. Quite a number of general-purpose programming languages with focus or at least direct support for XML data have been proposed recently, e.g., XMLambda [190], CDuce [29], XDuce [141], Xtatic (<http://www.cis.upenn.edu/~bcpierce/xtatic/>), Scriptol (<http://www.scriptol.com/>), Cw (<http://research.microsoft.com/Comega/>, [189]), and with special focus on Web services and composition XL [104, 105], Scala [199], Water [214]. All of these languages provide some form of specialized data structures for representing and accessing XML data. For existing programming languages, convenient access to XML data can be achieved using some API such as DOM¹, SAX², or XmlPull³ or by means of a language extension, e.g., HaXML [256] for Haskell, XMerL [258] for Erlang, or XJ [132] for Java.

However, when considering reasoning-aware query languages, the distinction between general-purpose programming languages and query languages becomes blurred, as such query languages are often computationally complete (cf. 5). For the purpose of this survey, a pragmatic approach has been chosen: a language is included, if querying is a core aspect of the language design or the approach to accessing Web data is unique and not covered by other proposals.

- *Evolution and reactivity*. A reactive system allows the specification of the dynamic aspects of a data storage system, i.e., (a) what changes are allowed (b) how to react when a certain event, such as the insertion or deletion of some data occurs. Several proposals for adopting previous approaches such as ECA rules to a Web setting have been published recently. Obviously, there is a close relation between languages for specifying the reactive behavior of a system and those for querying the current state as provided by conventional query languages: reactive languages often employ some query language for evaluating whether (a) the current event matches any of the reactive rules and (b) for conditional rules whether the data is currently in a status matching that condition. However, for this survey only reactive languages that also provide interesting querying abilities are considered. For a survey of reactive languages for the Web, refer to [10].
- *“Rule languages”*. Transformations, queries, derivations and reactive behavior can often conveniently expressed in rules. Recently, considerable interest in formalizing the rules guiding business decisions in such a way that they can be (a) understood and possibly even managed without learning a complicated rule syntax, (b) changed rapidly without refactoring existing programs, and (c) used directly to automate or support business

¹ <http://www.w3.org/DOM/> ² <http://www.saxproject.org/> ³ <http://www.xmlpull.org>

decisions such as whether a certain customer may receive a loan or which supplier to use for a certain part. This interest has also triggered the development of numerous, often proprietary languages for “rule engines”, i.e., systems that allow the specification and evaluation of such rules often as part of so-called expert systems. Examples for languages often used in this context include Prolog, F-Logic, and various extensions of these languages. In the Web context, the serialization and exchange of rules is particularly interesting as demonstrated by, e.g., the RuleML [44] initiative.

Again where to draw a line between query languages and general rule languages is not obvious. As in the cases above, in this survey only languages focusing on the efficient and effective querying of data are considered with exceptions for approaches that provide interesting insight for querying.

In the future, it might be interesting to extend the languages covered, e.g., to investigate differences and similarities with respect to requirements, principles, and realization of query languages and reactive, general rule, or programming languages in the Web context.

As stated in the introduction, this survey focuses on the language aspect of querying the Web. Therefore, (a) authoring tools such as visual editors are only considered in the context of the query language they are based upon and (b) issues related to storing and indexing Web data are not addressed (for a survey on storage system for RDF refer to [174]).

Despite all these conscious limitations in the kind of languages to be considered, the number of languages still remaining is still surprisingly large. This demonstrates the increasing interest in the area of Web query languages, in particular in RDF and Topic Maps query languages for which the respective standardization bodies have recently started the standardization process for (in the case of RDF, low-level) query languages. However, in neither case the aim to develop a common query language supporting the representation formalisms expected to be at the core of a future Semantic Web has received sufficient priority.

Following the overview of the languages considered, a concise summary of the evaluation results is given in Chapter 5, the details of which are given in Appendix B. From the evaluation results, a classification scheme for Web query languages is derived and briefly discussed by comparing it to previous approaches for classifying (Web) query languages and by demonstrating the ability to provide an insightful view on the languages surveyed here.

The paper is concluded by Chapter 6 with an outlook on possible improvements of this comparison and suggestions on interesting research directions derived from the evaluation.

Two appendices give (a) an overview of different serialization formats for RDF (Appendix A) and (b) the detailed results of the evaluation in tabular format (Appendix B).

Chapter 2

Preliminaries

In this chapter, a concise overview of the three representation formalisms that form the bases for the query languages investigated in this work is given. In particular, a small collection of sample data is described against which queries for assessing the functionality of the query languages considered are evaluated.

The “Extensible Markup Language (XML) is a simple, very flexible text format derived from SGML [...]. Originally designed to meet the challenges of large-scale electronic publishing, XML is also playing an increasingly important role in the exchange of a wide variety of data on the Web and elsewhere.”¹ XML [52] The previous quote hints at the dual role that XML is currently performing in the Web context: First and foremost, XML provides means for defining the syntax of new languages simplifying specification and deployment considerable, as common issues such as character encoding, markup syntax, linking (ID/IDREF, XLink [93]), mixing markup from different languages, splitting and reassembly of data fragments (XInclude [184]), etc. are handled uniformly at the level of XML.

However, this is not the only reason for using XML: More and more XML, and the underlying semi-structured data model, is recognized as a flexible means for representing, exchanging, and processing heterogeneous data originating from different sources. In this sense, an XML document can be interpreted as a rooted, directed, non-ranked, ordered graph. For many applications one does not consider the various reference or linking mechanisms defined for XML as part of the data model, thus reducing the interpretation of an XML document to a non-ranked, ordered tree. Although this is the data model adopted by the W3C (cf. XML Infoset [86] and XQuery 1.0 and XPath 2.0 data model [100]) and the majority of XML query languages, it is nevertheless recognized that providing means for traversing relations beyond the parent-child and sibling relation conveyed in the tree. Examples for such relations are links established by common keys in ID and IDREF attributes, XLink relations that can also be typed just as in RDF (cf. [87] for a more detailed analysis of the commonalities of RDF and XLink), or application-dependent relations.

As XML has been designed with focus on the first role, some peculiarities such as attributes or the lack of a standardized means, to express that the order of the children of some element is irrelevant and does not be preserved, make processing XML not always as convenient as one might hope, nevertheless XML is and most likely will remain the foundation for most Web application that require the exchange of data and increasingly also for applications where such

¹ <http://www.w3.org/XML/>

exchange is not needed.

Basically, an RDF [156, 27] model can be seen as an oriented graph whose nodes are labeled by either URIs, which describe (Web) resources, or literals (elementary data such as strings or numbers), or are unlabeled (the so called anonymous or “blank nodes”). The nodes are connected by arcs, also labeled by URIs, which are intended to represent “properties” of nodes (so blank nodes can be used to “aggregate” properties). A common alternative view of such a graph is a set of triples, called “statements”, of the form *(Subject, Property, Object)*, where *Subject* and *Object* are graph nodes, and *Property* is an arc. While this model accounts for the use of RDF as a very general description framework of Web resources, properties with special meanings are predefined in the RDF [135] and RDFS [135] specifications [179, 156, 135, 54] to describe, for instance, that a node is the “type” (`rdf:type`)² of another one, or is sub-class (`rdfs:subClassOf`) or sub-property (`rdfs:subPropertyOf`), etc. RDFS also defines a number of meta-classes, such as `rdfs:Class`, the class of all classes, or `rdfs:Property`, the class of all properties. The inheritance model exhibits some peculiarities, viz., (a) that resources can be classified in different classes that are not related in the subsumption hierarchy, (b) that the subsumption hierarchy can be cyclic (so that all classes on the cycle are equivalent), (c) that properties are first class objects and, in contrast to most object-oriented subsumption hierarchies, one does not describe which properties can be associated with a class but rather can specify the domain and range of properties. Based upon the information provided by an RDFS schema (or, to use another termed used almost equivalently in this case, ontology) certain inference rules can be specified, e.g., for inferring the transitive closure of the subsumption hierarchy or the type of an untyped resource that has a property associated for which the domain is known. OWL [188, 238, 24] extends the means provided by RDFS for defining the vocabulary used in describing resources.

RDF is designed for the exchange of meta-data represented as resource descriptions in an RDF graph. Therefore, a syntax for serializing and transferring RDF data is required. However, early approaches for an XML syntax of RDF have raised considerable critique, mostly for being overly complex to understand and process. Therefore, a large number of alternative serialization formats have been proposed. Appendix A presents a detailed overview of these serialization formats.

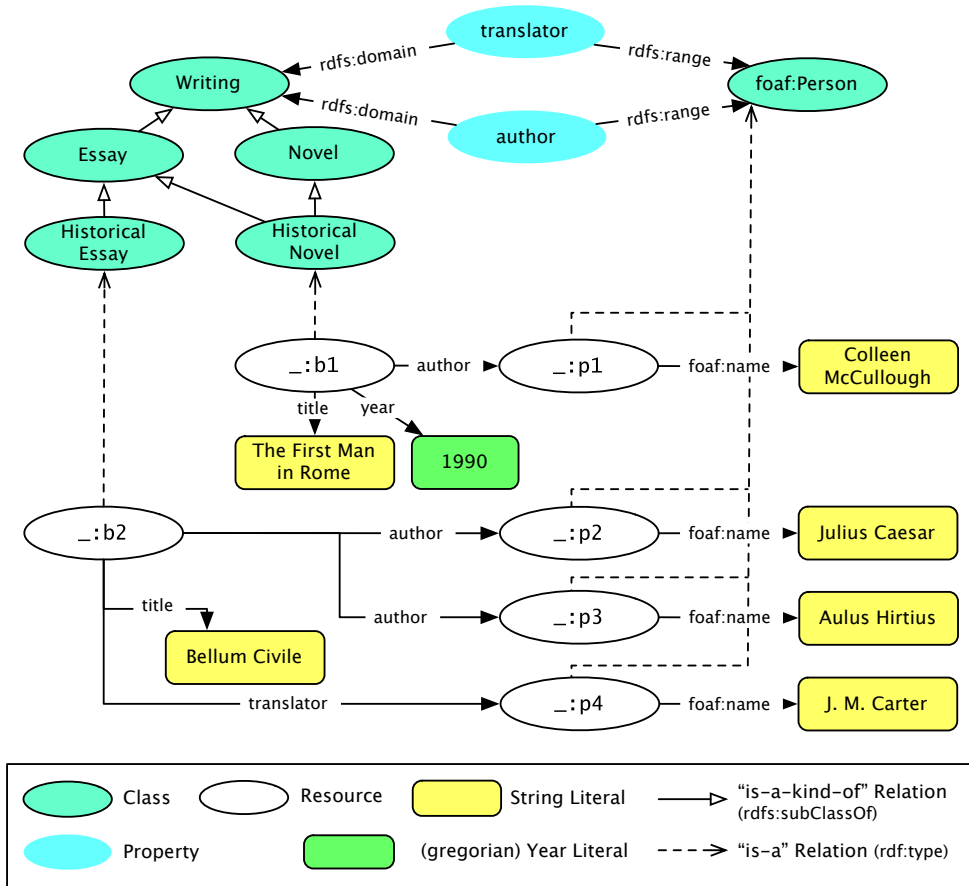
The last representation formalism that forms the basis of some of the query languages investigated in this survey is the ISO Topic Maps standard [144]. Inspired by previous work in the library sciences and on knowledge indexing, the Topic Maps standard [144] defines a data model consisting in a rich set of modeling primitives for representing, structuring, indexing, and relating knowledge. As the previously discussed representation formalisms at the core a topic map is a unranked graph with labeled edges and nodes. The most notable differences to RDF are the richer modeling primitives, that include, e.g., the ability to scope any information provided about a topic and to provide multiple facets of information. Also instead of binary associations, Topic Maps provides *n*-ary associations with roles for distinguishing the members of an association (rather similar to the XLink model for extended links). Topic Maps provides a basic ontology language for specifying a hierarchy of types of topics and associations.

The similarity of Topic Maps and RDF has been recognized and first efforts for integrating the two formalisms are presented in [160, 116].

From the perspective of this survey, all three representation formalisms can be used to

² Here and in the rest of this paper, common prefixes such as `rdf`, `rdfs`, `owl`, `xsd`, or `xtm` are assumed to be associated with the appropriate namespace.

Figure 2.1 Sample Data: Graphical representation of an RDF/S graph



represent the sample data discussed below varying mostly in the degree to which standard vocabulary is provided for defining the ontology part of the sample data.

2.1 Collection of Sample Data

For reasons of brevity and consistency, all queries operate on the same data, a collection of information on books. Figure 2.1 shows a graphical representation of an RDF/S graph. For more details see [156, 167, 54]. Note, that some of the RDF statements are represented in a more compact form, e.g., all resources with type `rdfs:Class` are depicted as special nodes instead of explicitly showing the `rdf:type` relation. Also, special arrows are used for the `rdf:type` and `rdfs:subClassOf` relations. Resources are identified with temporary IDs in the fashion of N3 [32], e.g., `_:b1`. Note also that the graphical notation used does not make explicit the connection between the property nodes `translator` and `author` (depicted by blue ellipses) and the instances of this property.

The sample data contains a small ontology using only the subsumption (or “is-a-kind-of”) relation `rdfs:subClassOf` and the instance (or “is-a”) relation `rdf:type`. This ontology is used to illustrate some of the specific requirements for a Semantic Web query language. We believe, that it is sufficient to show the most interesting issues involved in ontology querying without adding unnecessary complexity. Note, however that both the following discussion and the criteria for the evaluation are for general Web query languages including languages for the standard Web only, i.e., where such ontology information can not be represented in a standardized way but rather using an application dependent vocabulary. Furthermore, several aspects of the underlying data representation formalisms, such as RDF and Topic Maps, influence the desiderata for a Semantic Web query language even when no ontology is involved.

Since all three representation formalisms use the XML Schema simple datatypes defined in [38] for typing scalar data. The book titles and the names of the authors are string literals (either untyped or typed as `xsd:string`). The publication year of a book is typed as Gregorian year (`xsd:gYear`).

The sample data is assumed to be stored at the URL `http://example.org/books#3`. Where useful, this URL is associated with the prefix `books`, e.g., for referencing the vocabulary defined in the ontology part of the data.

Appropriate **(textual) representation of this data** in the different representation formalisms are chosen as basis for the actual queries. Note, that, since the issue of this survey is not to compare the different representation formalisms, we deliberately chose data and queries that can be handily represented in any of the formalisms.

Sample data in RDF. As the graphical representation is based on the RDF version of the data, this is shown first using the Turtle serialization syntax proposed in [26], a subset of N3 [32] (for more details on the syntax used, cf. Appendix A and the citations).

```

@prefix rdf:    <http://www.w3.org/1999/02/22-rdf-syntax-ns#> .
@prefix rdfs:  <http://www.w3.org/2000/01/rdf-schema#> .
@prefix xsd:   <http://www.w3.org/2001/XMLSchema#> .
@prefix foaf:  <http://xmlns.org/foaf/0.1/> .

:Writing a rdfs:Class ;
  rdfs:label "Novel" .
:Novel a rdfs:Class ;
  rdfs:label "Novel" ;
  rdfs:subClassOf :Writing .
:Essay a rdfs:Class ;
  rdfs:label "Essay" ;
  rdfs:subClassOf :Writing .
:Historical_Essay a rdfs:Class ;
  rdfs:label "Historical Essay" ;
  rdfs:subClassOf :Essay .
:Historical_Novel a rdfs:Class ;
  rdfs:label "Historical Novel" ;
  rdfs:subClassOf :Novel ;
  rdfs:subClassOf :Essay .
:author a rdfs:Property ;
  rdfs:domain :Writing ;
  rdfs:range foaf:Person .
:translator a rdfs:Property ;
  rdfs:domain :Writing ;
  rdfs:range foaf:Person .
_:b1 a :Historical_Novel ;

```

³ The URL is chosen in accordance to RFC 2606[97] on the use of URLs in sample data.

```

        :title "The First Man in Rome" ;
        :year "1990"^^xsd:gYear ;
        :author [foaf:name "Colleen McCullough"] .
_:b1    a :Historical_Essay ;
        :title "Bellum Civile" ;
        :author [foaf:name "Julius Caesar"] ;
        :author [foaf:name "Aulus Hirtius"] ;
        :translator [foaf:name "J. M. Carter"] .

```

The RDF serialization is, as expected, rather straightforward. Note that both the books and their authors and translators are represented by anonymous nodes (either without identifier or with a temporary identifier indicated by the `_:` prefix).

Sample data in Topic Maps. For the Topic Maps version of the data, the rather compact and readable Linear Topic Maps syntax [110] is used. The subclass-superclass associations are identified using the published subject identifiers defined in XTM [212]. For illustrative purposes the title of a book is represented as an occurrence of that topic. Finally, it is worth mentioning that this representation has been chosen more to demonstrate different features of the query languages surveyed than as a natural expression of the data in Topic Maps. One might, e.g., prefer to use a `publication` association that connects a book with its publisher, the year of publication, and the edition. Also, instead of separate associations for author and translator one could also provide a generic association between persons and books and use appropriate roles for differentiation.

```

/* Association and topic types for subclass-superclass hierarchy */
[superclass-subclass = "Superclass-Subclass Association Type"
 @ "http://www.topicmaps.org/xtm/1.0/core.xtm#superclass-subclass" ]
[superclass = "Superclass Role Type"
 @ "http://www.topicmaps.org/xtm/1.0/core.xtm#superclass" ]
[subclass = "Subclass Role Type"
 @ "http://www.topicmaps.org/xtm/1.0/core.xtm#subclass" ]
/* Topic types */
[Writing = "Writing Topic Type"
 @ "http://example.org/books#Writing" ]
[Novel = "Novel Topic Type"
 @ "http://example.org/books#Novel" ]
[Essay = "Essay Topic Type"
 @ "http://example.org/books#Essay" ]
[Historical_Essay = "Historical Essay Topic Type"
 @ "http://example.org/books#Historical_Essay" ]
[Historical_Novel = "Historical Novel Topic Type"
 @ "http://example.org/books#Historical_Novel" ]
[year = "Topic Type for a gregorian year following ISO 8601"
 @ "http://www.w3.org/2001/XMLSchema#gYear" ]
[Person = "Person Topic Type"
 @ "http://xmlns.org/foaf/0.1/Person"]
[Author
 @ "http://example.org/books#author" ]
[Translator
 @ "http://example.org/books#translator" ]
/* Associations among the topic types */
superclass-subclass(Writing: superclass, Novel: subclass)
superclass-subclass(Writing: superclass, Essay: subclass)
superclass-subclass(Novel: superclass, Historical_Novel: subclass)
superclass-subclass(Essay: superclass, Historical_Essay: subclass)
superclass-subclass(Essay: superclass, Historical_Novel: subclass)
superclass-subclass(Person: superclass, Author: subclass)

```



```

superclass-subclass(Person: superclass, Translator: subclass)
/* Occurrence types */
[title = "Occurrence Type for Titles"
  @ "http://example.org/books#title" ]
/* Association types */
[author-for-book = "Association Type associating authors to books"]
[translator-for-book = "Association Type associating translators to books"]
[publication-year-for-book = "Association Type associating translators to books"]
/* Topics, associations, and occurrences */
[p1: Person          = "Colleen McCullough"]
[p2: Person          = "Julius Caesar"]
[p3: Person          = "Aulus Hirtius"]
[p4: Person          = "J. M. Carter"]
[b1: Historical_Essay = "Topic representing the book 'The First Man in Rome'"]
author-for-book(b1, p1: author)
publication-year-for-book(b1, y1990)
{b1, title, [[The First Man in Rome]]}
[b2: Historical_Novel = "Topic representing the book 'Bellum Civile'"]
author-for-book(b2, p2: author)
author-for-book(b2, p3: author)
translator-for-book(b2, p4: translator)
{b2, title, [[Bellum Civile]]}

```

Sample data in XML. Here one of many possible XML representations of the sample format is shown. For brevity, the information that authors and translators are persons is not represented. Also, note the use of ID/IDREF links for representing the subsumption data.

```

<bookdata xmlns:xsd="http://www.w3.org/2001/XMLSchema#">
<book type="Historical_Novel">
  <title>The First Man in Rome</title>
  <year type="xsd:gYear">1990</year>
  <author>
    <name>Colleen McCullough</name>
  </author>
</book>
<book type="Historical_Essay">
  <title>Bellum Civile</title>
  <author>
    <name>Julius Caesar</name>
  </author>
  <author>
    <name>Aulus Hirtius</name>
  </author>
  <translator>
    <name>J. M. Carter</name>
  </translator>
</book>
<category id="Writing">
  <label>Writing</label>
  <category id="Novel">
    <label>Novel</label>
    <category id="Historical_Novel">
      <label>Historical Novel</label>
    </category>
  </category>
<category id="Essay">
  <label>Essay</label>
  <category id="Historical_Essay">
    <label>Historical Essay</label>
  </category>
  <category idref="Historical_Novel" />

```

```
</category>  
</category>  
</bookdata>
```

Alternatively, an XML serialization of the Topic Maps or RDF data shown above could be used. However, the serializations of both RDF and Topic Maps in XML are rather awkward and would only complicate the understanding of the issues involved in querying XML.

Chapter 3

Evaluation Criteria

This survey is guided by a number of evaluation criteria divided roughly in five areas: ease of use, functionality, semantics, formal properties and implementation, ontology awareness, and reasoning abilities. The evaluation criteria have been derived (1) from a number of relevant use cases and requirements specifications for web query languages [176, 203, 255, 81, 19, 115, 69, 68], (2) from a close look at the capabilities and intended scenarios for deployment for the query languages considered here, and (3) from the design principles and guidelines for the development of a standard and Semantic web query language described in [60].

As discussed in the introduction, this survey is focused on the suitability of the considered approaches for querying the Semantic Web. This emphasis is reflected in many of the following criteria, in particular in the areas of ontology awareness and reasoning abilities.

In this section, the evaluation criteria are introduced, described and an attempt is made to provide a justification for the selection of these evaluation criteria. In Chapter 4, a large number of XML and Semantic Web query languages are closely investigated along the criteria established here. From this investigation, one can easily observe that the proposed criteria allow the identification of several interesting classes of query languages reflecting different approaches, philosophies and requirements for querying the (Semantic) web.

3.1 Ease of Use

It has been previously mentioned that one of the most important aspects when designing a query language (or in fact any language to be used, at least partially, by human beings) is the “feeling” of the language or, in other words, how easy it is to use the language for a given (reasonable) task. Obviously, this is a highly subjective question and hard to measure without empiric studies. Therefore, in Chapter 4 the design philosophy and rational of several query languages are illustrated in more detail and a set of queries is used to give a concrete impression of the languages and their differences and similarities.

Aside of these restrictions, the following aspects of a query language are investigated to give some impression of how convenient the use of the language is:

Syntax. Query languages are often tailored to a specific perception as to who will author queries: whereas expert users usually prefer a *human-readable textual syntax* (C 1.1.), for be-

gainers or casual users even a simple textual syntax might already be too intimidating, however an appropriate *visual syntax* (C 1.2) or a *natural language interface* (C 1.3) can often make a language accessible to such users. Some query languages, e.g., [196], are designed for automatic generation by programs, therefore automatic query manipulation, e.g., by means of an *XML syntax* (C 1.4), is essential. Furthermore, in particular in the semantic web context, the automatic adaptation of queries, e.g., based on ontological data, is an essential issue. In such cases, *meta circularity* (C 1.5) is a desirable language feature, as it allows the adaptation of queries without involving learning and maintaining an additional programming language and environment.

- 1.1 *Human-readable textual syntax*. Does the query language provide a human-readable textual syntax.
- 1.2 *Visual syntax*. Is there a graphical editor or visual syntax for the query language?
- 1.3 *Natural language syntax or interface*. Is it possible to express queries in some kind of (usually restricted) natural language, e.g., in a variant of ACE [108].
- 1.4 *XML syntax*. Has an XML¹ syntax been specified for the query language.
- 1.5 *Meta circularity*. Is the query language capable of processing queries written in (at least) one of its syntactic forms.

Extensibility and Modularity. To be able to support a wide range of users with different degree of knowledge and expertise, leads to a set of interesting properties of a query language. Above the use of different syntactical representations of a query for different usage scenarios of the query language has been discussed. Furthermore, many languages provide some mechanism to allow queries to be written in a *modular* (C 1.6) fashion. E.g., views or rules can be provided by more experienced users, if necessary, and used by beginners transparently.

A similar aspect is that of *extensibility* (C 1.7): Different users requires often different functionalities and it is neither desirable, nor in all cases possible to provide all functionality within a single query language or processor. Furthermore, an extensible query language will be better equipped to adapt to emerging use cases in the future.

- 1.6 *Query modularity*. What kind of constructs for writing modular queries does the language support? Such constructs can be views (cf. [175]), rules (cf. [230, 237]), functions (cf. [41]), etc.
- 1.7 *Extensibility*. Does the language have a well-defined extension mechanism? Is it possible to detect from within the language what extensions are available in a given environment?

Adherence to Conventions. Both a shallow learning curve and the reuse of previously obtained expertise are supported if a query language uses established conventions where possible. Since this survey considers query languages not only for a single data representation formalism, such as XML, Topic Maps, and RDF, an important criterion for the evaluation is which *data representation formalisms* (C 1.8) are supported by a query language. Although

¹ For this survey, the only syntax considered particularly convenient for automatic processing of queries is XML, as this is the only syntax used for this particular purpose by the surveyed languages.

both OWL and RDFS are based on RDF, they are considered separate for this survey, as query languages are likely (and in fact do, cf. OWL-QL) to use the abstract model of OWL rather than query the concrete RDF serialization. Furthermore, both OWL and RDFS introduce constructs (e.g., `rdfs:subClassOf` or `owl:TransitiveProperty`) with a specific semantic that should be supported by a query language.

Related but separate from the issue of the supported data representation formalism is the *data model* (C 1.9) used by the query language. The use of a familiar and appropriate data model will certainly reduce the time a user requires to become acquainted with the query language. Although at first glance one might suspect that all languages for each of the above data representation formalisms use a uniform data model, this is not the case. E.g., some XML query languages consider the data to be strictly hierarchical, i.e., a tree, others offer support for ID/IDREF or similar linking mechanisms, hence use a graph data model. Similarly, some RDF query languages consider RDF data as mere triples, i.e., relational tuples with fixed arity 3 (e.g., [232, 194]), some as arbitrary graph (e.g., [148]), some restrict the graph to be acyclic or rooted (e.g., [208]).

Aside of the data model, also a *syntactical similarity* (C 1.10) with an existing language might be helpful for a novice user. However, often such similarities are merely superficial and can actually impede the understanding of a query language as the intuition from the existing language might not apply or at least not apply in all cases. Therefore, the question, on which *programming paradigm* (C 1.11) a query language is based upon, is often more illustrative of the abilities and general “feeling” of a query language.

Finally, all of the representation formalisms for the Web considered here are based on a graph or tree data model, therefore requiring some *accessor constructs* (C 1.12) that allow the access to specific nodes in the graph or tree based on their position (or relation) to other nodes. Query languages for structured data can be roughly classified by the accessor constructs they provide: *Pattern-based* query languages allow access to several parts of the graph at once specifying the relations among the accessed nodes by tree or graph patterns. *Path-based* query languages use constructs similar to file-system paths to access (usually) a single set of nodes based on any number of relations with other nodes in the graph specified in the path expression. Path-based languages can be further divided in languages that provide only true paths as accessors and languages where it is possible to describe tree-like queries (cf. [78]). Finally, *step-based* query languages provide only constructs for querying the relation of two (sets of) nodes. If it is to be queried whether more than a single relation holds for a certain set of nodes, the multiple relations have to be queried separately and joined via variables. A typical example for a step-based query language is RDQL [232], examples for path-based query languages are XPath [78], RQL [148], and RDFPath [208], Xcerpt [230] is a pattern-based query language.

1.8 *Data representation formalisms.* Which of the data representation formalisms (viz. XML, RDF, RDFS, OWL, Topic Maps) are supported by the query language?

1.9 *Data model.* What data model is used by the query language?

1.10 *Syntactical similarity.* Is there a strong syntactical similarity to other query languages, e.g., SQL, XQuery, or XPath?

1.11 *Programming paradigm.* What programming paradigm is the query language based upon?

1.12 *Accessor constructs.* Is the language based on single steps, paths, or patterns for specifying which nodes in a graph or tree structure to access.

3.2 Functionality

Complementary to the ease of using a query language is its functionality. This survey focuses on three aspects for measuring the functionality provided by a query language: What kind of queries can be expressed in the query language? For which of the concepts of the underlying representation formalism(s) are adequate query constructs provided? Are issues like updates, integrity constraints and active rules considered?

The emphasis on evolution and reactivity might be considered odd, and, indeed, almost none of the languages and systems analyzed in Chapter 4 does provide an update language, let alone means for specifying reactive behavior. Nevertheless, for the Web in general and even more for the Semantic Web, there is a clear need for sophisticated reactive components that allow the fast propagation of and reaction on changes in the (possibly remote) data and other events. Therefore, although support for evolution and reactivity is hardly a very discriminating criterion for the query languages considered in this survey, it is included to illustrate that a strong integration of query languages and reactive behavior is essential for the Semantic Web.

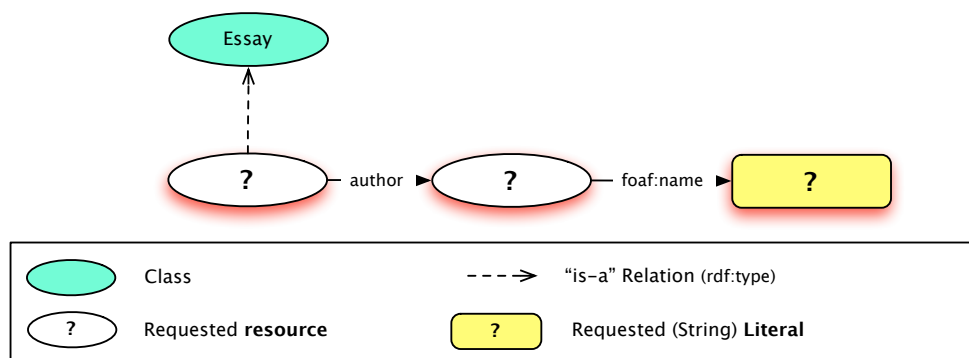
The following discussion details each of these points in order and proposes a set of evaluation criteria that are deemed to be useful for judging what functionality is provided by a query language.

3.2.1 Supported Query Types

For the purpose of this survey five classes of queries have been identified based upon previous work on classifying query languages by the provided functionality, most notably [176, 81]. To illustrate these queries and for later reference, some exemplary queries on the book data from Figure 2.1 are given both in natural language and in an easy-to-understand graphical notation based on the data graph.

- **Selection and Extraction Queries:** (C 2.1.1) The most basic type of query is to ask for some of the information represented in the data, usually based on its content, structure or position within the entire data.

Query 1. “Select all essays together with their authors and the names of their authors.”



Even such basic queries already raise a number of interesting issues with respect to the capabilities of a query language:

- *Supported result formats* (C 2.1.2). The query languages considered in this survey differ quite notably in this respect. One of the reasons is the different data representation: Is the data represented in XML, one could, e.g., return a set or sequence of all book elements (possibly already containing the authors and their names) or construct new elements grouping the authors under the related books. However some XML query languages, most notably XPath, do not support the construction of new elements, but always return a set or sequence of the selected elements. For RDF data, one might expect that the statements relating books and authors and the ones associating names to the authors will be returned. However, several query languages for RDF [LIST] do not return triples, but rather a table with one column per variable in the query and one row for each result. Similar considerations apply to Topic Maps.
- *Selection of related information* (C 2.1.3). As in the sample query, one is often not only interested in one piece of information (e.g., the books), but also in related information (e.g., the authors and their names). Usually, it should be possible to obtain the relations from the result, e.g. in the above case one would usually like to know which book relates to which author.

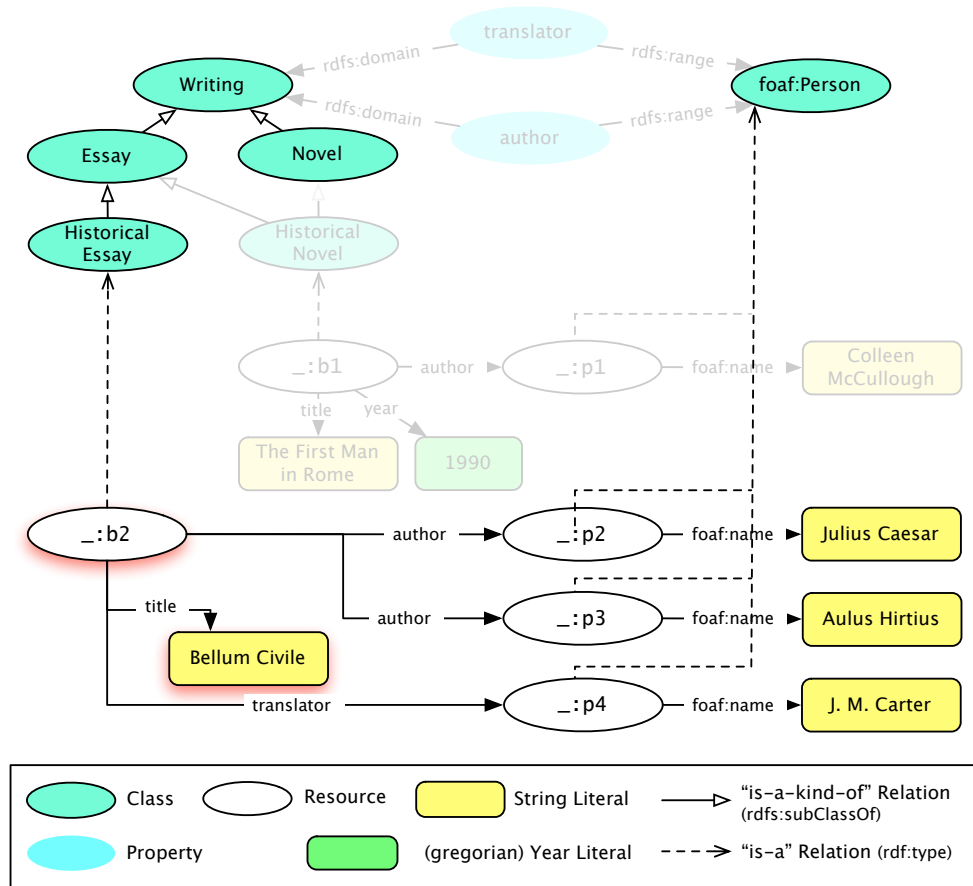
Again, there are a number of path-based query languages, such as XQL [227], XPath 1.0 [78] or RDFPath [208] that do not provide the ability to select related information.

But there are also cases where one is simply interested in the information itself without the relations among them. Since representing the relations among the information pieces in the result is expensive, it is desirable to allow both forms of returning result. Whereas XML query languages usually allow the author of a query to make this distinction, most RDF query languages do not provide this possibility.

- Should only books directly classified as “essay” be returned or also ones that are classified in one of its subclasses, e.g., as “historical essay”? In this case, the query actually involves inference, see below.

Another flavor of selection queries that is particularly relevant in a Semantic Web context for collection all information about a particular resource or topic are queries that *extract entire substructures* (C 2.1.4) from the data, e.g., a subgraph of an RDF graph.

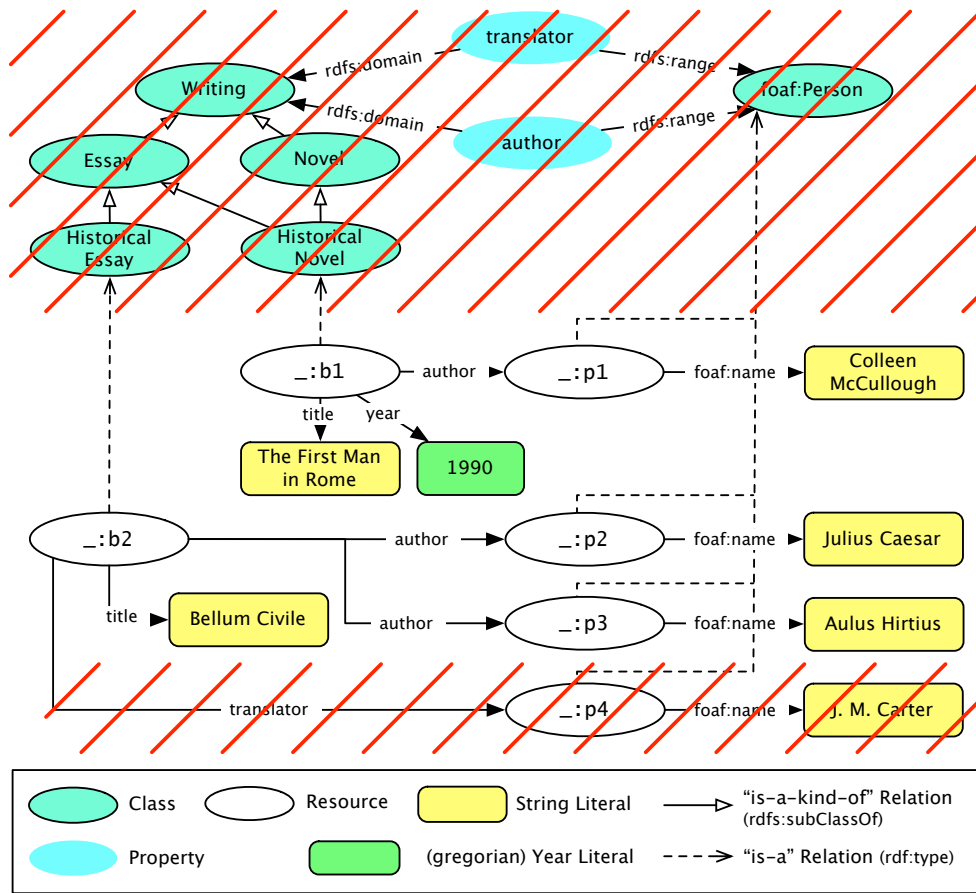
Query 2. “Select everything related to the book with title ‘Bellum Civile’.”



- **Reduction Queries:** (C 2.1.5) In some cases, instead of specifying precisely what to return, it is easier to specify what should *not* be returned as result. One might, e.g., not be interested in any ontology information or in translators of books for a certain application. The ability to specify what is *not* to be returned, is required, e.g., if the schema of the data to be retrieved is not known in advance, but the schema of the data to be left out is (at least to some extent).

Unless some specific support for reduction queries is offered by a query language, the specification of what should not be returned often requires some form of *negation* (C 2.1.6).

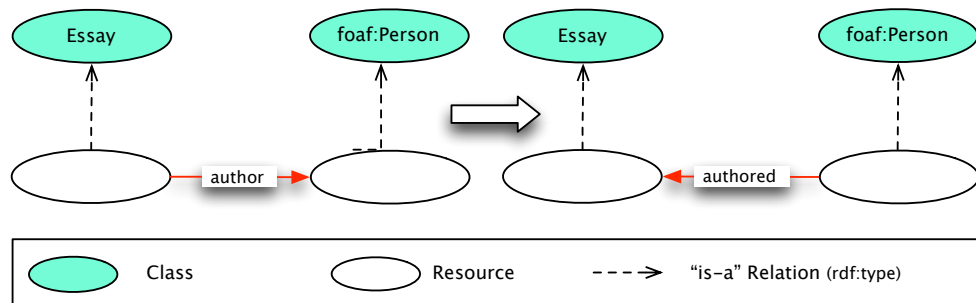
Query 3. "Select everything except all ontology information and any translators."



In the context of the Semantic Web, reduction queries become even more relevant, e.g., for combining information from different sources or for handling trust issues.

- **Restructuring Queries:** (C 2.1.7) Whenever structured data is to be queried, it is imperative to be able to change not only the value of the data but also its structure.

Query 4. "Invert the relation 'author' from a book to an author to 'authored'."



In RDF, restructuring of data is used to express statements about statements: A statement that is to be used as subject of another statement is “reified” by assigning an identifier *I* (i.e., an URI) to the statement and transforming the original statement into three statements for specifying subject, predicate and object of the statement.

- **Aggregation Queries:** A simple form of derivation of new knowledge, and hence linked to inference, is the aggregation of data. When structured information is considered, one can not only *aggregate on data values* (C 2.1.8) , cf. Query 5 but also on information about the *structure* (C 2.1.9), as shown in Query 6.

Query 5. “Return the latest year (the maximum of all years) in which an author with name ‘Julius Caesar’ published anything (i.e., any ‘Writing’).”

Query 6. “Return each of the subclasses of ‘Writing’ together with the average number of authors per publication of that type.”

Related to aggregation are the concept of *grouping* (C 2.1.10) and *sorting* (C 2.1.11) of the result returned. Note, that grouping and sorting is not meaningful for all of the representation formalisms that form the basis of the query languages discussed here, e.g., in RDF statements do not have any intrinsic order, however sequence container allow the specification of sequences.


- **Combination and Inference Queries:** Often it is necessary to *combine* (C 2.1.12) existing but not explicitly connected information, e.g., from different sources or represented in varying structures. Many ontologies specify, e.g., which names or identifiers are synonymous, i.e., refer to the same entity.

Query 7. “Combine all information about a book named ‘The Civil War’ and authored by ‘Julius Caesar’ with the information about the book with identifier `bellum_civile`.”

Combination of existing information often allows the *inference* (C 2.1.13) of additional data: if the two books named “Bellum Civile” and “The Civil War” are the same book and “Julius Caesar” is an other of “Bellum Civile” then he is also an other of “The Civil War”.


Another important form of inference queries are queries or views that compute the transitive *closure* of relations such as the `subClassOf` relation used in RDF for defining a subsumption hierarchy.

Query 8. “Return the transitive closure for `subClassOf` relation.”



Not all inference queries are combination queries, as the following example illustrates, where a new relation is (intensionally) defined based on existing data about books:

Query 9. “Return the co-author relation between two persons that stand in author relation with the same book.”



Whereas some query languages provide special closure operators for specifying which relations are transitive, others limit queries on the transitive closure to a fixed set of relations, e.g., the `subClassOf` relation from RDFS. Finally, some query languages provide a *general recursion* mechanism, that allows among others to query the transitive closure of arbitrary relations, even those defined intensionally as above.

- 2.1.1 *Selection queries.* Does the query language support selection queries?
- 2.1.2 *Supported result formats.* In what serialization or representation formalism can the result be returned? Possible values are, e.g., input subset (only a subset of the input can be returned), arbitrary XML, restricted XML (with an explanation of the restriction), one of the RDF serializations (cf. Appendix A, one of the Topic Map serializations, table (usually with one column for each variable to be included in the result and one row for each match), set or sequence (for query languages lacking the ability to select related information).
- 2.1.3 *Selection of related information.* Is it possible to return related information within the same query?
- 2.1.4 *Substructure extraction.* Are the means provided to select an entire substructure (e.g., a subtree or subgraph), in particular if the extent of the substructure is not known in advance?
- 2.1.5 *Reduction queries.* Does the query language support reduction queries?
- 2.1.6 *Negation.* Is it possible to express negation in the query language, e.g., to test the non-existence of data or to specify reduction queries.
- 2.1.7 *Restructuring queries.* Is it possible to rearrange the structure of the input or to create an entirely new structure?
- 2.1.8 *Aggregation on data values.* Can data values be aggregated in the query language?
- 2.1.9 *Aggregation on structure.* Is it possible to aggregate over the structure of the data, e.g., to determine the maximum number of authors for a book?
- 2.1.10 *Grouping.* Does the query language support grouping?
- 2.1.11 *Sorting.* Does the query language support sorting?
- 2.1.12 *Combination queries.* Is it possible to combine information that is not directly connected by the structure, e.g., by means of a join over some identifier?

- 2.1.13 *Inference queries*. What means for inference queries are provided by the query language? Possible values are, e.g., closure on predefined relations, closure on arbitrary structural relations (e.g., parent-child and sibling relation in XML or any property in an RDF graph), general recursion.

3.2.2 Adequacy

Considering the often still immature and evolving nature of the current representation formalisms for the Web and the often frustrating lack of a common and formal understanding of the underlying data models, it is not surprising that there is a great variability among the supported features and interpretations of the underlying representation formalisms among the query languages discussed here: Exemplary issues are, e.g., whether to consider an XML document with ID/IDREF links as graph or tree data or whether to support the more controversial features in RDF, e.g., reification and containers.

Therefore, we consider for each of the representation formalisms separately a number of criteria to give an impression of what a query language supports. These criteria are based on observing where the surveyed query languages differ and not meant to be a comprehensive list of features of the representation formalisms.

- **XML:** The XML data model (as defined in [86]) is an ordered tree and therefore provides two basic structural relations among nodes (representing elements in the document) in a tree: the *parent-child* (C 2.2.1) and the *sibling* (C 2.2.2) relation. XML query languages should support queries involving either relation and their *closures* (C 2.2.3), i.e., ancestor-descendant and preceding-following relations. Note, that XML languages that employ paths for accessing nodes often allow such relations to be queried in both directions, e.g., from the parent to the children as well as from a child to its parent. However, in [207] it has been shown that a restriction to queries where the relations are followed in “forward” direction (w.r.t. the order in which nodes are visited) is reasonable. A third relation particular to the XML data model is provided via *ID/IDREF links* (C 2.2.4). If these relations are handled transparently, an XML document actually has to be represented as a graph (cf. C 1.8).

Support for the intrinsic order of an XML document often goes beyond merely providing constructs for querying the sibling relation: many XML query languages allow to *access nodes by their position* (C 2.2.5) within some sequence of related nodes, e.g., the second title of each book. Some query languages also allow *unordered matching* (C 2.2.6), where the order among siblings is ignored. Finally, all XML query languages *preserve the order* (C 2.2.7) when returning some part of the input unless the result is reordered specifically.

One particular important aspect when considering XML data is that the data often lacks a fixed schema or the schema allows for a certain amount of flexibility. Therefore any XML query language should be able to specify both *partial queries* (C 2.2.8), i.e., queries where only some constraints on the data are specified and the existence of additional nodes in the data does not affect the matching, and *total queries* (C 2.2.9), i.e., queries that only match if there is no more data than specified in the query. Also often desirable is the ability to specify *optional* (C 2.2.10) parts of a query: if there is some data matching an optional sub-query it will be returned, but if no data matches the query other related items are still returned: E.g., one might want to select all books together with their trans-

lators but still return the books if there is no translator for it. Such optional sub-queries resemble outer or inner joins in relational databases.

Another important distinction among XML query languages is the support for *construction of new elements* (C 2.2.11). It is prerequisite for supporting restructuring queries.

Finally, there are some issues concerning the alignment of an XML query language with the emerging standards defined for XML: support for *XML Schema* [98, 250, 38] (C 2.2.12)—more details on typing are discussed below in Section 3.6.1—, support for *namespaces* (C 2.2.13), support for advanced linking using *XLink* [93] and *XPointer* [125, 124, 91, 92] (C 2.2.14), and construction of compound documents specified with *XInclude* [184] (C 2.2.15).

- **RDF:** Beyond simple triple statements that associate two resources (called subject and object) via a certain property (or predicate), the RDF data model has some peculiarities that require special attention when discussing a query language for RDF. Even more than in the case of XML, RDF has been design under the assumption that in a global Semantic Web fixed schema information is often unattainable. Therefore, all properties in RDF are *optional* (C 2.2.10) and *multi-valued* (C 2.3.1), e.g., a book described in RDF can have any number of authors including none. A query language should provide suitable constructs for optional sub-queries as in the case of XML. Furthermore, properties are, as any resource, identified by an URI and can therefore be the subject of other statements. In this case, it is required that a query language is able to *query (the identifiers) of properties* (C 2.3.2), and not only those of subjects and objects. Another aspect of this are containers and collections in RDF: *containers* (C 2.3.3) provide the means for expression sets, sequences and alternatives of resources, e.g., to express that a committee has voted in a certain way without implying that every single member actually voted in that way. Within sequences *access by position* (C 2.2.5) is often useful, e.g., to obtain the first in the sequence of chapters for a book. Alternatives also have to be handled differently from sets or sequences since it can only be derived that at least one of the objects is actually related to the subject (by the given predicate). *Collections* (C 2.3.4), introduced only recently during the revision [27] of the original RDF proposal [167], differ from containers in their semantics: collections can be “closed” in the sense that no further elements can be included in the collection (for example when consulting additional descriptions on the same resource). This is not the case for RDF containers. Note, that both containers and collections can be reduced to triples (i.e., binary relations), however requiring considerable effort by the user. Therefore, specific support for these constructs takes a considerable burden from the query programmer.

A similar consideration applies to one of the more controversial abilities of RDF, viz. to be able to express statements about statements via *reification* (C 2.3.5). As stated above, A statement that is to be used as subject of another statement is “reified” by assigning an identifier *I* (i.e., an URI) to the statement and transforming the original statement into three statements for specifying subject, predicate and object of the statement. An RDF query language should provide at least support for the transparent access to reified statements, i.e., the query programmer should not have to specify whether a statement is reified or not.

E.g., for the attribution of what the source of a statement is, a statement is often associated with additional information about its context. Although such context informations

are often useful and desirable [127, 154], this can only be realized in RDF using reification. Therefore, the Jena Toolkit [126], one of the syntactical forms for RDF [229] discussed in Appendix A and a recent query language [220] all use the concept of “*quads*” (C 2.3.6), i.e., triples enriched with an additional context or source information. This allows not only a more compact formulation of queries but can also be used for improved storage of such statements (cf. [95]).

As mentioned in Chapter 2 and discussed in further detail in Appendix A, there is a plethora of different serialization formats for RDF. Therefore, it is interesting to note, which *serializations* (C 2.3.7) can be used for input and for the result of a query.

Issues related to typing and classification of resources are considered when discussing RDFS, OWL and the ontology awareness

- **Topic Maps:** Topic Maps [144], as discussed, are an ISO standard for representing “information about the structure of information resources”. As such, they have several commonalities with RDF, but are separated in a clearer focus on the application area in which they are intended to be used. In [19] and [115] use cases and requirements for a upcoming Topic Map query language have been detailed. A large number of these requirements are covered by criteria discussed in other sections of this survey. However, from previous approaches for the conceptualization of information, Topic Maps inherit a number of rich modeling concepts that should be supported by a Topic Maps query language: In contrast to the binary relations of RDF, Topic Maps provide *n-ary associations* (C 2.4.1) with *labeled roles* (C 2.4.2) identifying the associated data. As in RDF, association can be further described by other associations, requiring that the query language supports the *querying of association names* (C 2.3.2). Also similar to RDF and XML, which associations can be used for a certain topic type is not fixed, thus *optional* sub-queries (C 2.2.10) are often helpful. There is a number of *predefined associations* (C 2.4.3), e.g., for defining a sub-class hierarchy similar to RDFS, some of them with special semantics (e.g., transitivity). Also query languages should support some of the more advanced concepts of Topic Maps, such as *scopes* (C 2.4.4), that allow to limit the validity of an association and *facets* (C 2.4.5) used to create filters.

Within Topic Maps, subject identity can be established by the use of subject indicators (also known subject descriptors). A query language should be able to connect information on the same topic provided, e.g., from different sources, by *querying these subject indicators* (C 2.4.6).

As with RDF, there is a number of different proposals for *serializing Topic Maps* (C 2.4.7), e.g., XTM [212], LTM [110], and AsTMa= [16], and it should be noted which of these are supported by a Topic Maps query language.

One of the most prominent differences between Topic Maps and RDF is that essentially in Topic Maps all “statements” are already reified, i.e., for each association there is a node representing that association (for more details, [244]. Therefore, no special treatment of reified “statements” is necessary in Topic Maps.

Aside of the concrete aspects of the different representation formalisms, there is a number of features for a query language that are provided by languages specifically designed to be used in a Web context. Such query languages have to deal with inherent heterogeneity in a global distributed system. *Querying multiple data sources* (C 2.5.1) and providing *multiple*

(often different) answers to multiple outputs (C 2.5.2) is just the basis to such querying. The query language should also provide means to reduce the amount of data being transferred between Web sites, e.g. by *limiting the result size* (C 2.5.3), support for *reduction queries* (C 2.1.5) limiting the answers to what is actually of interest to the user and by supporting for easy *distribution of sub-queries* (C 2.5.4) based on the data sources accessed, e.g., by clearly identifying which parts of a query deal with what data. This relates to the *composability of queries* (C 2.5.5): Often the result of one query is to be further refined, possibly at another Web site. To maximize interoperability, the query language should support different *output serializations* (C 2.5.6) and provide some mechanism for the *user to specify* (C 2.5.7) which serialization to choose.

Since data sources are heterogeneous (e.g., w.r.t. the level of structure in the data provided) and may, in general, contain erroneous or contradicting data, means for *vague or approximate query answering* (C 2.6.1) and *relevance ranking* (C 2.6.2) of query answers are often necessary to deal with Web data sources. Approximation should be possible not only as in classical information retrieval on the text content of the data but also on its structure. To some extent, this is already covered by the above discussion, e.g., through constructs expressing optional sub-queries that allow for certain variation in the data to be matched. However, some query languages might provide additional constructs for *approximate structure matching* (C 2.6.3). Due to the varying level of structure provided by the data sources, some kind of text processing can often not be avoided. Therefore, Web query languages might also include classical information retrieval features as proposed for querying XML data, e.g., in [109] and in the recent extensions of XQuery and XPath [12]. Such features include *full-text queries* (C 2.6.4) such as single word-search, phrase search, ordered multi-word search, and proximity search based on word distance or structure, and *word normalization* (C 2.6.5), e.g., stemming, stop-word handling, suffix and prefix removal, or the use of thesauri, dictionaries or taxonomies.

Data sources are heterogeneous not only with respect to the data representation, but also in their communication abilities. E.g., some data sources might provide data rather slowly, therefore making it undesirable to wait with processing that data until all necessary data has actually arrived. In such a case, an implementation providing *progressive processing* (C 4.12) of the data can often provide answers long before all data has arrived. The dual case are data sources that provide data so rapidly that conventional techniques for parsing and storing that data (either in memory or in a database) are infeasible. A *streamed* (C 4.11) implementation of a query language evaluates queries directly against the stream of incoming data without requiring expensive data structures to be built.

- 2.2.1 ^{xml} *Parent-child relation*. Does the query language provide access to the parent-child relation?
- 2.2.2 ^{xml} *Sibling relation*. Is it possible to query the siblings of a node?
- 2.2.3 ^{xml} *Closure relation*. Are there means to query the closure of the two base relations parent-child and sibling.
- 2.2.4 ^{xml} *ID/IDREF*. Does the query language support the explicit or transparent dereferencing of ID/IDREF (or similar) linking mechanisms? Possible values for this criterion are: transparent (i.e., links are automatically resolved and can be queries like parent-child relations), explicit (i.e., there is some specific construct to be used for following ID/IDREF links), indirectly (i.e., it is possible to dereference ID/IDREF links but only by querying the actual XML attributes—no specific constructs are provided) and none.

- 2.2.5 ^{xml, rdf} *Access by position.* Is it possible to access nodes in the structure based upon their absolute or relative position within the structure? In RDF, this only applies for (sequence) containers.
- 2.2.6 ^{xml} *Unordered matching.* Is it possible to specify several children for the same node in the XML structure such that the order among the children is not relevant for finding a match?
- 2.2.7 ^{xml} *Order-preserving result.* If the result of a query is a part of the input, is it possible to preserve the order?
- 2.2.8 ^{xml} *Partial queries.* Is it possible to specify only partial constraints on the data to be matched by a query, e.g., allowing for additional children of a node to exist?
- 2.2.9 ^{xml} *Total queries.* Is it possible to specify that a match must fulfill exactly the constraint given in the query, disallowing the existence of additional data.
- 2.2.10 ^{xml, rdf, tm} *Optionality.* Is it possible to express that certain sub-queries are optional, i.e., that their result should be included in the total result if they match anything, but if they do not match, the remainder of the query may still yield result?
- 2.2.11 ^{xml} *Construction.* Does the query language provide means for constructing new elements and attributes?
- 2.2.12 ^{xml} *XML Schema.* Does the query language make use of XML Schema [98, 250, 38] or a similar schema language such as RELAX NG [79, 80], e.g., for typed queries or type checking of programs or validation of result?
- 2.2.13 ^{xml} *Namespaces.* Does the query language support XML namespaces as defined in [50, 51]?
- 2.2.14 ^{xml} *XLink and XPointer.* Is there some specific support for extended linking mechanisms as provided by XLink [93] and XPointer [125, 124, 91, 92]?
- 2.2.15 ^{xml} *XInclude.* Does the query language allow the construction of compound documents using, e.g., XInclude[184]?
- 2.3.1 ^{rdf} *Multi-valued properties.* Is it possible to query and return properties with multiple values?
- 2.3.2 ^{rdf, tm} *Querying property or association identifiers..* Is it possible to query the identifiers of properties or associations, e.g., for finding properties of properties?
- 2.3.3 ^{rdf} *Containers.* Does the query language have some provisions for container support, e.g., for construction of containers, querying a sequence?
- 2.3.4 ^{rdf} *Collections.* Are there some specific constructs for querying and returning collections (defined in [135])?
- 2.3.5 ^{rdf} *Reification.* Is there some specific support for reification in the query language? Such support can be, e.g., the transparent querying of reified statements or an easy notion for querying and constructing reified statements.

- 2.3.6 ^{rdf} *Quads*. Does the query language offer specific support for context or source information associated with a statement (often represented as “quads” [229])?
- 2.3.7 ^{rdf} *RDF serializations*. Which of the different RDF serializations discussed in Appendix A are supported by the query language and its implementations?
- 2.4.1 tm *Querying n-ary associations*. Is it possible to query relations of arbitrary arity?
- 2.4.2 tm *Labeled roles*. Is it possible to use roles for querying and construction and to query the role labels?
- 2.4.3 tm *Predefined associations*. Does the query language support predefined associations such as subclass-superclass relationship and their special semantics?
- 2.4.4 tm *Scopes*. Is their explicit support for scopes in the query language?
- 2.4.5 tm *Facets*. Is their explicit support for facets in the query language?
- 2.4.6 tm *Subject indicators*. Are (published) subject indicators supported for connecting information from different data sources?
- 2.4.7 tm *Topic Maps serializations*. Which of the Topic Maps serialization formats are supported by the query language and its implementations?
- 2.5.1 *Multiple data sources*. Is it possible to query multiple data sources?
- 2.5.2 *Multiple outputs*. Is it possible to generate different answers for sending to multiple Web sites?
- 2.5.3 *Limiting result size*. Is there some mechanism for limiting the size of the result retrieved (possibly in conjunction with an SQL-like offset construct for paged answer retrieval)?
- 2.5.4 *Sub-query distribution*. Does the query language provide convenient means for identifying reasonable sub-queries for distribution to the data sources, if these provide such a capability?
- 2.5.5 *Composability of queries*. Is it possible to compose queries?
- 2.5.6 *Output serializations*. Which output serializations are supported by the query language? Here, only general values such as RDF, XML or Topic Maps are given. The different serializations for RDF and Topic Maps are investigated as criterion C 2.34 and C 2.41.
- 2.5.7 *User-specified serializations*. If there are different serializations provided, can the user choose the serialization (either implicitly by constructing the appropriate serialization directly or by selecting explicitly the desired output serialization)?
- 2.6.1 *Approximate query answering*. Is there a provision in the query language for answering vague or approximate queries?
- 2.6.2 *Relevance ranking*. Does the query language automatically rank answers by their relevance or provide some means for explicit ranking or scoring of answers?
- 2.6.3 *Approximate structure matching*. Does the query language support approximate matching only on the content or also on the structure of the data?

- 2.6.4 *Full-text queries.* Are there means for processing full-text content of the structural data to support, e.g., word or phrase queries?
- 2.6.5 *Word normalization.* If there is some support for full-text and word queries, is it possible to apply some word normalization such as stemming or normalization based on thesauri before matching?

3.2.3 Evolution and Reactivity

As discussed above, evolution and reactivity are important concepts linked to a query language that become even more important in the context of the Semantic Web. Three aspects of evolution and reactivity are particularly related to query languages (for a more detailed survey, cf. [10]):

- 2.7.1 *Update language.* Does the language provide updates or has a related update language been defined?
- 2.7.2 *Integrity constraints.* Is there some mechanism provided to define and enforce rules that specify restrictions on how the data can be changed?
- 2.7.3 *Event handling.* Is it possible to specify that certain actions should be performed if an event, such as an update, occurs?

3.3 Semantics

A clear and *well-understood formal semantics* (C 3.1) enables not only a better understanding of the workings of a query language, but also a certain independence from the actual implementation as queries written for one implementation of the language should be usable with any other implementation that follows the semantics. Furthermore, a formal semantics also proved to be very fertile for the development of various kinds of implementation-independent optimizations ranging from source-to-source transformations to operator reordering in logical query plans. These optimizations have, in contrast to optimizations on the level of the physical query plan, the advantage that they usually are usually not specific to a single implementations, but rather can be applied to a number of implementations (e.g., with similar characteristics). Such optimizations often benefit from two characteristics of a formal semantics: *compositionality* (C 3.2) and *referential transparency* (C 3.3), as these enable “local” optimizations where the context in which a sub-query occurs does not have to be considered. Finally, for any kind of reliable reasoning (e.g., where also proof traces are to be delivered), a formal and well-understood semantics is indispensable.

As a reference for implementations and as basis for cost estimations used in query optimization, an *operational semantics* (C 3.4) with a rigid mapping into the formal semantics is often desirable. Such an operational semantics can, e.g., be provided by means of an *abstract machine* (C 3.5) together with a translation from the formal semantics into instructions of that abstract machine.

- 3.1 *Formal semantics.* Has a (well-understood) formal semantics been established for the query language?

- 3.2 *Compositionality*. Is the semantics of the language compositional, i.e., defined in such a way that the semantics of a compound construct is based on the semantics of its parts?
- 3.3 *Referential transparency*. Is the semantics of the query language constructs referentially transparent, i.e., not depending on the context?
- 3.4 *Operational semantics*. Has an operational semantics (together with a mapping from the formal into the operational semantics) been defined?
- 3.5 *Abstract machine*. Has an abstract machine for the query language been defined?

3.4 Formal Properties and Implementation

Aside of a formal semantics, also the status and properties of implementations for a query language is worth noting. This gives some impression on the capabilities of the query language that are also related to formal properties such as complexity and completeness.

In the past, in particular for relational databases, two properties have often be touted to be distinguished features of query languages in contrast to general purpose programming languages: *declarativeness* (C 4.1), i.e., that a query language describes what should be result and not how that result can be obtained, and no *computational (Turing) completeness* (C 4.2).

[228] defines a declarative language as a language where each operation is declarative, i.e., “*independent* (does not depend on any execution state outside of itself), *stateless* (has no internal execution state that is remembered between calls), and *deterministic* (always gives the same results when given the same arguments)” [228]. This definition is adopted for the purpose of this survey.

Historically the lack of computational completeness found in many relational query languages such SQL has been perceived as an advantage both for query writing and query execution and optimization. Where computational completeness has been required either stored procedures written in some general-purpose programming language or the embedding of a query language into a host programming language have been used. However, in the last decade it has been recognized (e.g., during the development of the latest SQL standard—sometimes collectively referred to as “SQL3”) that computational completeness actually provides benefits in many cases: Embedding in a host language and stored procedures have both proven to be troublesome with respect to interoperability, performance (in particular for embedding in a host language) and authoring of queries. In this survey, we therefore also note whether a query language is *computationally complete* or not.

In particular for expressive (or even computationally complete) query languages, the issue of scalability w.r.t. efficiency is crucial: It should be possible to characterize *interesting sub-languages* (C 4.3) that can be evaluated with different complexity, such that simple and frequent queries can be evaluated rather quickly, but queries using more expressive features of the query language might actually take longer. In particular, it should be noted, whether a *polynomial core* (C 4.4) has been identified for the query language, i.e., a sub-language such that all queries written in that language can be evaluated in polynomial (combined) complexity. Often the evaluation strategies for such sub-languages differ, therefore an efficient *automatic classification* (C 4.5) of queries in terms of their computational cost is highly desirable.

Additional to the formal properties discussed, this survey also includes a brief overview over the **implementations** provided for a query language. The following issues are investi-

gated: *number* (C 4.6) of different implementations, *status* (C 4.7) of the implementation (e.g., prototype, internal production use, external production use), and *support* (C 4.8) by major database or Web technology vendors such as Oracle, IBM, HP, Microsoft, or large open-source projects such as Apache.

As discussed above, different kinds of implementations of a query language are very desirable in a Web context with varying application requirements and capabilities of data sources. Therefore we note for each query language whether it has been implemented on top of a *database* (C 4.9) (i.e., for querying persistent data where updates are rare and queries are frequent), in an *in-memory query processor* (C 4.10) (i.e., where both query and data are used once only), and in a *streamed* (C 4.11) fashion (i.e., where persistent or continuous queries are evaluated against volatile data). In each of these cases, answers might be provided once all data has been processed or in a *progressive* (C 4.12) manner, i.e., as soon as possible.

- 4.1 *Declarativeness*. Can the query language be considered declarative following the definition from [228]?
- 4.2 *Computational completeness*. Is the query language computationally complete?
- 4.3 *Interesting sub-languages*. Have interesting sub-languages been defined, e.g., with different complexity characteristics?
- 4.4 *Polynomial core*. Has a sub-language with polynomial complexity been proposed?
- 4.5 *Automatic classification*. If there are sub-languages with different complexity characteristics, has a method been proposed to automatically classify queries?
- 4.6 *Number of implementations*. How many different implementations have been developed for the query language?
- 4.7 *Status of implementations*. What is the status of these implementations?
- 4.8 *Support*. Is there support for the query language by major database or Web technology vendors or large open-source projects.
- 4.9 *Database implementation*. Has the query language been implemented on top of a database?
- 4.10 *In-memory implementation*. Has an in-memory processor for the query language been developed?
- 4.11 *Streamed implementation*. Has a streamed implementation been provided or considered?
- 4.12 *Progressive implementation*. Has a progressive implementation been provided or considered?

3.5 Reasoning

Reasoning or the ability to derive additional data based upon the actual data stored in the database has been an important ability of deductive or logic databases [252, 67] and reflected in query languages such as Datalog. Since reasoning is to be considered a defining element

in the Semantic Web vision, it is suitable to ask what kind of reasoning mechanisms query languages to be used in this context provide.

The first step to reasoning support is the ability to *intentionally specify* (C 5.1) data, e.g., by means of rules, views or functions. Although this has already been considered under C 1.6 ('query modularity'), it will be noted here again to illustrate where simple derivations of new data are possible. To specify such a derivation, boolean operators or equivalent constructs are often used: *conjunction* (C 5.2, realized, e.g., by set intersection), *disjunction* (C 5.3, realized, e.g., by set union), *negation* (C 5.4, realized, e.g., by set difference), and *quantification* (C 5.5, realized, e.g., by relational division). Also being able to specify optional sub-queries (as discussed in C 2.2.10 ('optionality')) can ease the specification of derivations.

Based upon derivations as described in the previous paragraph, more powerful reasoning abilities can be provided. Recursion, in particular, allows the specification of complex derivations such as, e.g., the transitive closure of relations or associations. Several forms of recursions are provided by query languages: *general recursion* (C 5.6) where rules, views, functions or similar intentional data specifications can be recursive, *structural recursion* (C 5.7) where some means (e.g., a special operator for computing the transitive closure of a relation) for recursion along the structure of the data is possible, and transitive closure only over some *predefined relations with special semantics* (C 5.8).

Together with how to specify inference, it is also necessary to note what *inference theory* (C 5.9) is used by a query language (in other words, what kind of reasoning is actually provided). Since the Web is constantly evolving and the appropriate reasoning may differ between different domains, some query languages provide *theory extensibility* (C 5.10), i.e., provide a well-defined interface for adding new reasoners that either implement a different inference theory or provide additional reasoning abilities for specific domains (e.g., ontology reasoning by a description logics reasoner or temporal reasoning).

- 5.1 *Intensional data specification.* Does the query language support the intensional specification of data, e.g., by means of views?
- 5.2 *Support for conjunctions.* Does the query language support the use of conjunctions or similar operations?
- 5.3 *Support for disjunctions.* Does the query language support the use of disjunctions or similar operations?
- 5.4 *Support for negation.* Does the query language support the use of negations or similar operations?
- 5.5 *Support for quantification.* Does the query language support the use of quantification or similar operations?
- 5.6 *General recursion.* Is it possible to use general recursion in the query language?
- 5.7 *Structural recursion.* Is it possible to traverse the structure of the data recursively (e.g., by means of a recursive relation descendant in XPath [78])?
- 5.8 *Closure on predefined relations.* Are there some predefined relations that can be traversed recursively for accessing the transitive closure (e.g., `rdfs:subClassOf` in RQL [148])?
- 5.9 *Inference theory.* What inference theory is used by the query language, if any?

- 5.10 *Theory extensibility*. Is it possible to use different inference theories or to add reasoning abilities for specific domains?

3.6 Ontology Awareness

A query language to be used in a Semantic Web context should be able to incorporate ontologies: Ontologies can be used to query across data sources with different vocabularies for describing the data by providing a mediation between these vocabularies. They can also help to improve the recall of a query by extending the actual queried terms with semantic information (e.g., related terms, contextual information).

To leverage ontologies for querying, the ontologies and the semantic relations described within have to be queried. Although recent ontology languages such as RDFS [54] and OWL [188, 238, 24, 210] are building upon RDF (in that they provide a defined set of terms for RDF that allows the definition of new vocabularies), merely support for RDF is not sufficient to be able to access the knowledge contained in the ontologies appropriately. Rather, a query language has to be aware of the semantics of the terms provided by RDFS or OWL, e.g., the “is-a” relation (`subClassOf` relation in RDFS). Thus, the query language can use the ontology to derive new knowledge about the described instances, e.g., through property propagation. Therefore, the querying the *data described by the ontology and the ontology* (C 6.1.1) itself should be possible within the same language. The following issues are investigated to classify the level of support for ontology languages (for this purpose, only RDFS and OWL are considered, since these are the considered by some of the surveyed query languages):

- **RDFS:** RDFS [54] provides only a small set of terms for describing vocabularies in RDF. The semantics of these terms is defined in [135]. Some of these terms have specific properties, e.g., `rdfs:subClassOf` and `rdfs:subPropertyOf` are both transitive, others allow a limited form of reasoning, e.g., if a class C is the range (i.e., the set of possible values) of some property p , written as an RDF triple $(p, rdfs:range, C)$, and (x, p, y) (some x has the property p with value y , then one can infer that y is an element of class C , i.e., $(y, rdf:type, C)$).

The query languages surveyed here show different *support for RDFS* (C 6.2.1): Most languages treat RDFS terms as any other RDF term, i.e., without special consideration. Some languages support querying the transitive closure (C 6.2.2) of `rdfs:subClassOf` and `rdfs:subPropertyOf`, of which some provide transparent support, others require the explicit specification of the transitive closure (e.g., by means of a recursive rule, view or function or using a special closure operator). Only a small number of languages also use *RDFS for typing* (C 6.2.3), thus providing special constructs, e.g., to query the extent of a class, and static type checking.

- **OWL:** Based on previous ontology languages such as DAML+OIL, the recently specified Web Ontology Language (OWL) [188, 238, 24, 210] is starting to see wide-spread acceptance in academia and also for certain industrial applications. OWL supports a much larger set of terms for defining and constraining vocabularies. Since the support for OWL is still rather limited (with the notable exception of OWL-QL [103]), this survey only addresses some general issues related to *support for OWL* (C 6.3.1): support for *special property classes* (C 6.3.2) such as `owl:TransitiveProperty` or `owl:SymmetricProperty`

(e.g., when querying properties that are classified as `owl:TransitiveProperty` the transitive closure of the property should be used), *information propagation* (C 6.3.3) for classes in the subsumption hierarchy (defined by `rdfs:subClassOf`) and for class equivalence, intersection, etc., and *information propagation* (C 6.3.4) for individuals (e.g., when using `owl:sameAs`).

As in the case of RDFS, also OWL can be used for *typing* (C 6.3.5) the described data.

- 6.1.1 *Querying both ontology data and instance data.* Does the query language support querying the ontology data together with the data described by the ontology?
- 6.2.1 ^{rdfs} *Support for RDFS.* Does the query language support RDFS?
- 6.2.2 ^{rdfs} *Transitive closure for subsumption hierarchy.* Is it possible to query the transitive closure for the RDFS subsumption hierarchy (created using `rdfs:subClassOf` and `rdfs:subPropertyOf`)?
- 6.2.3 ^{rdfs} *Use of RDFS for typing.* Is the type system of the query language (at least partially based) on RDFS and the `rdf:type` relation?
- 6.3.1 ^{owl} *Support for OWL.* Does the query language support OWL?
- 6.3.2 ^{owl} *Special property classes.* Are the classifications of properties w.r.t., e.g., transitivity and symmetry, used for querying?
- 6.3.3 ^{owl} *Information propagation for classes.* Is information about classes (based on the subsumption hierarchy and class equivalence etc.) propagated for querying?
- 6.3.4 ^{owl} *Information propagation for individuals.* Is information about individuals propagated for querying?
- 6.3.5 ^{owl} *Use of OWL for typing.* Is the type system of the query language (at least partially based) on OWL and the `rdf:type` relation?

3.6.1 Type system

Related to ontologies and schema languages is the issue of **typing** in Web query languages. For the purpose of this survey, we have already introduced two criteria related to typing, viz. the use of ontologies specified in RDFS or OWL for typing information. For XML, type information can obviously also be provided by XML schema languages such as XML Schema [98, 250, 38] or RELAX NG [79, 80]. This survey only gives a brief overview of typing related questions in query languages (following [64]):

- 6.4.1 *Typing.* Have typing issues at all been considered for the query language?
- 6.4.2 *Static vs. dynamic typing.* Does the query language support static or dynamic typing.
- 6.4.3 *Explicit vs. implicit typing.* Is data typed by explicit type declarations or implicitly (e.g., by type inference as in the statically-typed Haskell or by type inspection as in the dynamically-typed Smalltalk).
- 6.4.4 *Type inference.* Does the query language provide type inference, e.g., to avoid explicit type declarations.

6.4.5 *Type coercion.* Is it possible to change the type of an expression, either automatically or using, e.g., a cast operator?

6.4.6 *Support for XML Schema simple data types.* Both RDF and XML Schema use the same set of “simple” data types defined in [38]. Therefore, this criterion notes whether these simple data types are supported by the query language?

Based upon these 111 evaluation criteria, the following section presents an overview of the evaluation followed by a short description of the evaluated languages.

Chapter 4

Query Languages for the Web: An Overview

4.1 XML Query Languages

In this chapter, a very brief overview of some XML query languages is provided. For reasons of space and time, only some of the vast number of XML query languages proposed are considered. In particular, no visual query are left out for they are hard to evaluate by the evaluation criteria discussed here. Recently, there has been a number of proposals for combining full-text querying capabilities provided in information retrieval systems with structured access to XML data. Again, for reasons of space and time, these languages have not been considered.

4.2 Textual XML Query Languages

4.2.1 Navigational Languages

4.2.1.1 Lorel

Lorel [5] is a query language originally designed for semistructured data (more specifically, the language OEM [209, 118]) that was later adapted to XML data. Its syntax strongly resembles SQL and OQL, but it is capable of navigating graph structures in a path-like fashion.

Query 10 (Lorel). Select all authors and titles of books written after 1991 and return them in `result` elements contained within a `results` element.

```
select xml(results:
  (select xml(result:
    (select X.author, X.title
     from bib.book X
     where X.@year > 1991))
  ))
```

Lorel is a navigational language with rule-like `select-from-where` queries. It is capable of querying several documents and evaluating joins, but order is not considered when querying, as semistructured data is usually always unordered. Multiple data items can be

retrieved by assembling several path selections. Although the result of a query is a set of object identifiers (OIDs), XML elements can be constructed using expressions of the form `xml(tagname:subexpression)`. As it descends from OQL, it is capable of grouping and supports aggregations. Queries can be nested, in which case query and construction are not separated. Lorel supports a basic type system but is not aware of XML schema information and does not use type inference.

4.2.1.2 XPath

The *XML Path Language* (XPath) [78] is a W3C recommendation for a selection language whose primary purpose is to address parts of an XML document. Since it lacks construction and reassembling aspects, it cannot be considered a full-fledged query language and is thus called a *selection language* in this thesis. Many other XML query languages build upon XPath, most prominently XQuery and XSLT.

XPath expressions specify navigation steps within the data tree represented by a document, relative to a so-called context node (which is initially the root node of the document). An XPath expression consists of several location steps separated by `/`, each specifying *how to reach* a node relative to the previous node's position. XPath can thus be considered as a *navigational language*.

Query 11 (XPath). The following XPath expression selects titles of books with author Dan Suciu and year attribute with a value greater than 1991:

```
/bib/book[@year > 1991][author = "Dan Suciu"]/title
```

Literally, this expression reads: from the root, go to `bib` elements, from there to `book` elements for which holds that the attribute `year` is greater than 1991 and that contain an `author` element as child node with a text value of `Dan Suciu`, and select all `title` child nodes.

XPath differentiates axes like `child`, `descendant`, `parent`, `ancestor` or `sibling`. These axes can be classified into *forward axes*, which contain all such navigations that only move forward in the document tree and *backward axes*, which contain all such navigations that only move backward in the document tree [207]. Since XPath allows both forward and backward axes, evaluation can be very complex and a node might be visited several times during a selection. However, [207] shows that arbitrary XPath expressions without variables can be rewritten into equivalent XPath expressions containing only forward axes.

Forward XPath expressions resemble positional pattern, as they no longer specify arbitrary navigations through the document tree, with the minor exception that a pattern usually does not allow to match the same node in the document by two different nodes of the pattern:

Query 12 (XPath Forward). In the forward XPath expression

```
/a[child::b]/*[att="some value"]
```

the selection `*` and the `child::b` might match the same node in the database:

```
<a>  
<b att="some value"/>  
</a>
```

4.2.1.3 XQL

XQL [227, 225] is a variant of XPath that has been proposed and implemented by Microsoft and others in lieu of XPath becoming a completed recommendation. It differs from XPath only in minor points that are not relevant to this comparison.

4.2.1.4 XSLT

XSLT, the *Extensible Stylesheet Language* [77], is a language for transforming XML documents. Originally intended as a powerful style-sheet language, it is often considered as a query language as well, and the existence and development of two independent W3C XML query languages is often criticized. As it was the first available query language for XML, XSLT is very widespread and understood by many programmers. A multitude of implementations exist (e.g. as part of a standard library for XML processing in Java).

An XSLT style-sheet is composed of one or more transformation rules (called *templates*) that recursively operate on a single input document. Transformation rules are guarded by patterns, which are expressed in terms of XPath expressions. The first rule whose pattern matches is evaluated, all other rules are ignored. In contrast to most other query languages, XSLT uses an XML-only syntax:

Query 13 (XSL). Select all authors and titles of books written after 1991 and return them in result elements contained within a results element.

```
<xsl:stylesheet version="1.0"
                xmlns:xsl="http://www.w3.org/1999/XSL/Transform">
  <xsl:template match="/bib">
    <results>
      <xsl:apply-templates/>
    </results>
  </xsl:template>
  <xsl:template match="book[@year > 1991]">
    <result>
      <xsl:apply-templates select="title"/>
      <xsl:apply-templates select="author"/>
    </result>
  </xsl:template>
  <xsl:template match="title|author">
    <xsl:copy-of select="."/>
  </xsl:template>
</xsl:stylesheet>
```

This stylesheet is interpreted as follows:

- try to match the root node with the templates in the style-sheets (only first template matches)
- create a <results> element and within it try to recursively apply the templates to all child nodes of <bib>
- for each child node, if the year is greater than 1991, create a <result> element and recursively apply the rules to the <title> and author children of the context node
- for each <title> and <author> element, copy the complete input to the result.

Like XQuery, XSLT is based on the selection language XPath and thus a navigational language. Since a template always operates on a single node, neither retrieval of multiple data items nor joins are directly supported. It is, however, possible to assemble several XPath expressions within a construction pattern. XSLT always operates on a single input document and is not capable of retrieving data from more than one resource. It partly supports ordered/unordered data by using XPath. XSLT allows to construct new data and grouping, but aggregations are limited to those supported by XPath. Although XSLT is a rule-based language, it does not really support the separation of construction and querying, as each rule only applies to a single node. Rules may be called explicitly to form sub-queries, and such calls may be recursive (i.e. XSLT is Turing complete [151]). XSLT is an untyped language that is unaware of any available schema or type information.

4.2.1.5 XQuery/Quilt

XQuery [41] can be considered the “state-of-the-art” XML Query language as it is the current W3C recommendation for XML querying and therefore very widespread. XQuery has several predecessors, of which it resembles most the language *Quilt* [70], but influences from other languages (like XQL and XML-QL) are reflected in many constructs.

XQuery queries consist of so called FLWOR (FOR-LET-WHERE-ORDER BY-RETURN) expressions and use XPath (described above) for the selection of data items. FOR and LET serve to bind variables to values selected by XPath expressions. Whereas LET binds a variable to a set of data items, FOR iterates over the different data items in a set. The WHERE part may be used to specify conditions for the selected data items. ORDER BY has only been introduced recently and allows to order the results in a certain sequence. RETURN marks the beginning of a result pattern, which may contain additional XPath selections. XQuery expressions are enclosed in curly braces and embedded in the construction pattern.

Query 14 (XQuery). Select all authors and titles of books written after 1991:

```
<results>
{
  FOR $book in document("bib.xml")//book
  WHERE $book/@year > 1992
  RETURN <result>
        { $book/title }
        { $book/author }
  </result>
}
</results>
```

This query iteratively binds the variable \$book to all book elements occurring in the document in the FOR part. The WHERE part ensures that only such books are selected that have an attribute year with a value larger than 1991. The RETURN part gives a construction pattern that itself again contains subqueries for selecting the title and authors of the book.

Being based on XPath, XQuery classifies as a navigational query language. Multiple data items can be selected only by using multiple XPath expressions. XQuery supports arbitrary nesting of queries as well as the definition of external functions that may contain frequently used subqueries. XQuery is aware of both schema information and basic types and some implementations support static type checking (e.g. Galax¹). XQuery is not rule-based and heavily mixes querying and construction, making more complex queries difficult to read.

¹ <http://db.be11-labs.com/galax/>

4.2.1.6 FXT

The language *fx* [31], the *functional XML transformer*, is a transformation language which is similar to XSLT in that it uses the same kind of pattern-guarded rules to recurse over the input document. However, *fx* aims at optimal performance and thus puts certain limits on patterns and path expressions. Most importantly, *fx* neither supports explicit calling of rules (thus no recursion other than over the document tree) nor iteration constructs like XSLT's *for-each*. However, allows to perform auxiliary computations by embedding *SML* expressions in transformation rules.

Query 15 (fx). Select all authors and titles of books written after 1991 and return them in *result* elements contained within a *results* element.

```
<fxt:spec>
  <fxt:pat>/bib</fxt:pat>
  <results>
    <fxt:apply/>
  </result>
  <fxt:pat>//book</fxt:pat>
  <fxt:if test='fromString(Vector2String(getAttribute (String2Vector "year") current))
    > 1991'>
    <result>
      <fxt:apply select="/author"/>
      <fxt:apply select="/title"/>
    </result>
  </fxt:if>
  <fxt:pat>//author</fxt:pat>
  <fxt:copyContent/>
  <fxt:pat>//title</fxt:pat>
  <fxt:copyContent/>
</fxt:spec>
```

fx is based on the path language *fxgrep* and thus classifies as a navigational language. Interestingly, it allows to select (at most) two data items at once by using so-called *binary patterns*. The restricted path language allows neither joins nor to differentiate between ordered and unordered queries, and a transformation always operates on a single input document. *fx* allows to construct new elements, but aggregations and grouping can only be performed by reverting to the underlying functional language *SML*. *fx* is rule-based but neither allows rule chaining nor separates construction from querying. It supports basic types but does not take advantage of schema information and performs no type inference outside *SML* expressions.

4.2.1.7 XPathLog

LoPiX [186] is an implementation of the XML querying and data manipulation language XPathLog [187]. XPathLog aims at integrating F-Logic with path-based selection in XML documents. Queries in XPathLog are specified as conjunctions of path expressions in which variables may occur at multiple positions. It is thus possible to select several nodes in a single selection step. Further queries may refer to variable bindings. Two elements can be merged by so-called *object fusion*, which appears to be a mechanism similar to feature unification. This mechanism can also be used to group elements.

Instead of the transformation approach taken by most other query languages, where an input document is transformed into a new output document, XPathLog allows to update existing documents. Updates are specified by Prolog-like rules where the right-hand part consists of a query and the left-hand part specifies how to modify the document.

4.2.1.8 CXQuery

CXQuery, the *Constraint XML Query Language* [73], is an effort to create a declarative query language with support for schema information. It uses rules similar to Datalog and can use XPath to navigate the document tree as well as term-based patterns, but apparently, no deep structures are possible (as in Datalog). Construction is specified by term structures in the head of a rule. Again, nesting does not appear to be possible.

4.2.2 Positional Languages

4.2.2.1 XML-QL

XML-QL [94] is a positional, rule-based query language for XML and was designed at AT&T Labs. It uses an XML-based pattern language where variables may occur at arbitrary positions. An XML-QL query consists of a single CONSTRUCT-WHERE rule which may be divided into several subqueries.

Query 16 (XML-QL). Select all authors and titles of books written after 1991 and return them in result elements contained within a results element.

```
WHERE
  <bib>
    $book
  </> IN "bib.xml"
CONSTRUCT <results>
  WHERE <book year=$y>
    <title>$t</>
    <author>$a</>
    </book> IN $book, y > 1991
  CONSTRUCT <result>
    <title>$t</>
    WHERE $a2 IN $a
    CONSTRUCT <author>$a</>
  </result>
</results>
```

XML-QL patterns are positional, are capable of retrieving multiple data items at once from multiple sources and evaluate joins over variable name equality. XML-QL does not differentiate between ordered and unordered queries (everything is matched in any order). XML-QL can construct new data, but grouping can only be achieved by using nested sub-queries (as in the example above). XML-QL supports aggregations. Queries are rule-based, and sub-queries can be nested into the query rules (as above). Construction and querying are separated, but this separation is abandoned when using nested sub-queries. XML-QL is not aware of type information.

4.2.2.2 UnQL

UnQL [61], the *Unstructured Query Language* is a query language originally developed for querying semistructured data. UnQL uses a positional, pattern-based selection and a query consists of a single `select ...where ...` rule which separates construction from querying. Queries may also be nested, in which case the separation of querying and construction is abandoned. UnQL uses its own, non-XML syntax for representing and querying graph-structured data.

Query 17 (XML-QL). Select all authors and titles of books written after 1991 and return them in `result` elements contained within a `results` element.

```

select { results: (
  select { result: { title: T,
                    ( select { author: A }
                      where { author: A } in Book )
                  }
  }
  where { book: Book } in Bib,
        { year: Y, title: T } in Book ) },
        Y > 1991
where { bib: Bib } in db

```

In terms of properties, UnQL is very similar to XML-QL: it uses positional patterns, can retrieve multiple data items at once and from multiple sources, joins are possible over variable name equality. However, UnQL can respect the order of data, if desired. Like in XML-QL, construction is possible, but grouping can only be achieved by using nested sub-queries. Aggregations are supported. Queries are rule-based, and sub-queries can be nested into the query rules (as in the example above). Construction and querying are separated, but this separation is abandoned when using nested sub-queries. Like XML-QL, UnQL is not aware of type information.

4.2.2.3 XML-RL

XML-RL [170] is a proposal for a pattern-based query language based on logic programming. Patterns are expressed by terms that may contain logic variables and may be partly abbreviated with a path syntax similar to XPath. An XML-RL query program consists of several rules denoted by $A \Leftarrow L_1, \dots, L_n$ where A is used for construction and L_1, \dots, L_n are query patterns. Rules may interact via rule chaining and it is possible to use recursion.

Query 18 (XML-RL). Select all authors and titles of books written after 1991 and return them in `result` elements contained within a `results` element.

```

(file:result.xml)
/results/result: (title: $t, {author: $a})
←
(file:bib.xml)
/bib/book: (@year: $y, title: $t, author: $a), $y > 1991

```

XML-RL is a positional language, but allows path expressions as abbreviations. It can retrieve multiple data items from multiple sources and supports joins via variable name equality, as XML-QL and UnQL. It is not possible to query based on order. XML-RL can construct new data items. Grouping is possible by means of a special construct `{.}` and aggregations are supported. Queries programs can be structured by using more than one rule, and rules may interact via rule chaining. Querying and construction are always separated. However, computations are usually performed in the query part. XML-RL does not support type information.

4.2.2.4 XMAS

XMAS [172] part of MIX [21], the *XML Matching And Structuring language*, is a declarative, rule-based query language for XML. Rules are of the form `CONSTRUCT ...WHERE ...` and resemble XML-QL rules very closely. However, XMAS provides more powerful constructs for grouping/collection and aggregation, in a similar way to the grouping construct `{.}` of XML-RL.

Query 19 (XMAS). Select all authors and titles of books written after 1991 and return them in result elements contained within a results element.

```
CONSTRUCT
  <results>
    <result>
      $T
      $A {$A}
    </result> {$T}
  </results>
WHERE
  <bib>
    <book year=$Y>
      $T: <title/>
      $A: <author/>
    </>
  </> IN "bib.xml"
  AND $Y > 1991
```

Note the grouping expressed by enclosing the collected variables in curly braces. For every instance of \$T, a result element is created. Within every such result element, all authors are collected.

Like XML-QL, XMAS employs positional, rule-based selection, can query multiple data items at once and from multiple sources. In contrast to XML-QL, joins have to be expressed using an explicit join operator. XMAS is not capable of querying the order of elements, and does not support incompleteness in depth. New data can be constructed in rule heads using sophisticated grouping constructs that avoid nested subqueries. Apparently, aggregations and subqueries are not supported. Construction and querying is therefore always separated. XMAS is not aware of type information.

4.2.2.5 XET/XDD

XET [13], *XML Equivalent Transformations*, is a pattern-based, rule-based query language for XML aiming primarily at Semantic Web applications, but also capable of querying generic XML data. XET employs XML-based patterns enriched with variables for retrieving data items in XML documents. XET rules are similar to rules in logic programming and support rule chaining. A formal semantics is provided in form of XDD (*XML Declarative Description*) and ET (*Equivalent Transformations*). Unfortunately, there is not enough information available to express the running example in XET.

4.2.2.6 Xcerpt

Xcerpt [230] is a pattern-based query language for XML inspired by logic programming. It is further extended in the REVERSE I4 working group following the design principles detailed in [60]. An Xcerpt program consists of at least one *goal* and some (possibly zero) *rules*. Rules and goals contain query and construction patterns, called *terms*. Terms represent tree-like (or graph-like) structures. The children of a node may either be *ordered*, i.e. the order of occurrence is relevant (e.g. in an XML document representing a book), or *unordered*, i.e. the order of occurrence is irrelevant and may be chosen by the storage system (as is common in database systems). In the term syntax, an *ordered term specification* is denoted by square brackets [], an *unordered term specification* by curly braces { }.

Likewise, terms may use *partial term specifications* for representing incomplete query patterns and *total term specifications* for representing complete query patterns (or data items). A term t using a partial term specification for its subterms matches with all such terms that (1) contain matching subterms for all subterms of t and that (2) might contain further subterms without corresponding subterms in t . Partial term specification is denoted by *double* square brackets `[[]]` or curly braces `{ { } }`. In contrast, a term t using a total term specification does not match with terms that contain additional subterms without corresponding subterms in t . Total term specification is expressed using *single* square brackets `[]` or curly braces `{ }`.

Data Terms represent XML documents and the data items of a semistructured database, and may thus only contain total term specifications (i.e. single square brackets or curly braces). They are similar to *ground* functional programming expressions and logical atoms. A *database* is a (multi-)set of data terms (e.g. the Web). A non-XML syntax has been chosen for Xcerpt to improve readability, but there is a one-to-one correspondence between an XML document and a data term.

Query Terms are (possibly incomplete) patterns matched against Web resources represented by data terms. They are similar to the latter, but may contain *partial* as well as *total* term specifications, are augmented by *variables* for selecting data items, possibly with *variable restrictions* using the \rightsquigarrow construct (read *as*), which restricts the admissible bindings to those subterms that are matched by the restriction pattern, and may contain additional query constructs like *position matching* (keyword `position`), *subterm negation* (keyword `without`), *optional subterm specification* (keyword `optional`), and *descendant* (keyword `desc`).

Query terms are “matched” with data or construct terms by a non-standard unification method called *simulation unification* that is based on a relation called *simulation*. In contrast to Robinson’s unification (as e.g. used in Prolog), simulation unification is capable of determining substitutions also for incomplete and unordered query terms. Since incompleteness usually allows many different alternative bindings for the variables, the result of simulation unification is not only a single substitution, but a (finite) *set of substitutions*, each of which yielding ground instances of the unified terms such that the one ground term matches with the other.

Construct Terms serve to reassemble variables (the bindings of which are specified in query terms) so as to construct new data terms. Again, they are similar to the latter, but augmented by *variables* (acting as place holders for data selected in a query) and the *grouping construct* `all` (which serves to collect all instances that result from different variable bindings). Occurrences of `all` may be accompanied by an optional sorting specification.

Example. *Left:* A query term retrieving departure and arrival stations for a train in the train document. Partial term specifications (partial curly braces) are used since the train document might contain additional information irrelevant to the query. *Right:* A construct term creating a summarised representation of trains grouped inside a `trains` term. Note the use of the `all` construct to collect all instances of the `train` subterm that can be created from substitutions in the substitution set resulting from the query on the left.

<pre> travel {{ train {{ departure {{ station { var From } }}, arrival {{ station { var To } }} }} }}</pre>	<pre> trains { all train { from { var From }, to { var To } } }</pre>
---	---

Xcerpt also provides means for defining (possibly recursive) rules. Based upon these ability, an library of view definitions for convenient querying of RDF (and Topic Maps) is currently in development.

4.3 RDF Query Languages

4.3.1 SquishQL-family

4.3.1.1 SquishQL

SquishQL [194, 193] is an RDF query language that has been developed with ease-of-use and similarity to SQL as main principles. It has been implemented in at least three different processors, most notably in a slightly modified and further refined version, named RDQL, in the Jena Toolkit [126].

Essentially, SquishQL aims at a more intuitive, easy-to-use and fast-to-learn means for accessing RDF triples than provided by general-purpose RDF APIs. This aim also leads to a certain restriction with respect to the supported features.

The SquishQL query model is influenced by [129] and uses so-called “triple patterns” and conjunctions of these “triple patterns” to specify the structure of the RDF graph to be matched by the query. As stated in [194] “this results in quite a weak pattern language but it does ensure that in a result all variables are bound.”

To give an impression of the syntax used by SquishQL, the following example shows Query 1 formulated in SquishQL:

```
SELECT ?essay, ?author, ?authorName
FROM   http://example.org/books
WHERE  (?essay, <rdf:type>, <books:Essay>),
       (?essay, <books:author>, ?author),
       (?author, <books:name>, ?authorName)
USING  books FOR http://example.org/books#,
       rdf FOR http://www.w3.org/1999/02/22-rdf-syntax-ns#
```

Since RDQL is based on and mostly identical to SquishQL, refer to Section 4.3.1.3 for more information on this language family.

Project page:

- (Inkling: <http://swordfish.rdfweb.org/rdfquery/>, RDFStore: <http://rdfstore.sourceforge.net/>)

Implementation(s):

Inkling [193], Jena Toolkit—RDQL [233, 232, 231], RDFStore [222]

Online demonstration:

<http://demo.aseantics.com/rdfstore/www2003/>

4.3.1.2 rdfDB Query Language

rdfDB [128] is an early proposals for an RDF data base that influenced the design of, among others, SquishQL and RDQL (cf. Section 4.3.1.3). [96] gives an introduction into rdfDB, more details on the query language can be found in [128].

The syntax of rdfDB is SQL-like and as in SQL several different database commands can be executed in a single session, although no transactions management is provided. Aside of commands for creating databases, inserting and deleting triples, and defining namespaces, the core element of the rdfDB syntax is the `select-from-where` clause. As all database commands such a clause is delimited by a `</>` and returns bindings for any number of variables specified after the `select` such that the pattern specified in the `where` clause matches in the database given in the `from` part.

For illustration of rdfDB consider again Query 1 and a possible implementation of that query in the rdfDB query language (assuming the content of `http://example.org/books` has been stored in a database called `booksdata`):

```
enter namespace xmlns:books http://example.org/books# </>
enter namespace xmlns:rdf http://www.w3.org/1999/02/22-rdf-syntax-ns# </>
select ?essay, ?author, ?authorName from booksdata
where (rdf:type ?essay books:Essay), (books:author ?essay ?author),
      (books:name ?author ?authorName) </>
```

rdfDB is considered influential on the design of SquishQL and RDQL discussed in this section, but the development seems to have been halted. Nevertheless, an early implementation of rdfDB is still available from the project page.

Project page:

<http://www.guha.com/rdfdb/>

Implementation(s):

prototype available from the project page

Online demonstration:

none

4.3.1.3 RDQL

RDQL, RDF Data Query Language, is a query language for RDF models developed by Andy Seaborne at HP and recently submitted to the W3C as candidate for standardization [232, 194, 233, 231]. It is an evolution from several languages, inspired by the work in [129]: rdfDB [128], SquishQL [194], and Inking [193].

As SquishQL RDQL does not support the special meaning provided by the RDF schema level, although at least one of its implementations (viz. in the Jena Toolkit [126]) provide a transparent transitive closure over the subsumption hierarchies defined in RDFS using `rdfs:subClassOf` or `rdfs:subPropertyOf`.

A typical RDQL query has a syntax which is reminiscent of SQL, but which is built around a set of conjuncted triple patterns, i.e. triples of constants or variables, which are resolved on the graph. Such a process binds the variables to node or property labels, which are URIs and constants, and the result of the query is a subset of those bindings. Additional constraints on variable values can be used to filter the result. For instance, the query:

```
SELECT ?person
FROM   http://somewhere.org/some-rdf-model-of-people
WHERE  (?person, <http://example.org/peopleInfo#age>, ?age)
AND    ?age > 24
```

returns the URIs of persons whose age is greater than 24.

SELECT is the only language statement, and its syntax is easily described:

```
SELECT variables (identifies the variables whose bindings are returned)
FROM model URI
WHERE list of triple patterns
AND boolean expression (the filter to be applied to the result)
USING name FOR uri, ...
```

The last clause allows the simplification of queries by introducing names for long URIs. For instance, the previous query can be rewritten as:

```
SELECT ?person
FROM http://somewhere.org/some-rdf-model-of-people
WHERE (?person, info:age, ?age)
AND ?age > 24
USING info FOR <http://example.org/peopleInfo#>
```

The language is maintained intentionally simple, operating only on the “data” level of RDF, so that it could be easily amenable to standardization as a “low-level” RDF language, which relies on higher level services to make use of rules or inference facilities. As the author explicitly states, “if a graph implementation provides inferencing to appear as ‘virtual triples’ (i.e. triples that appear in the graph but are not in the ground facts), then an RDQL will include those triples as possible matches in triple patterns”, so that no distinction is made between inferred triples and ground triples [232].

This language philosophy, moreover, has the effect that the predefined properties which describe semantic oriented or “schema” aspects of a model, like type, set or class relations, are treated as ordinary properties, leading to cumbersome complex queries when those aspects are involved (like, for instance, a query to return all the elements of a container). This problem, too, could be solved by using other specialized wrappers around a model.

In the following, we will consider the main aspects of the language according to the classification criteria previously shown.

Easy of use The queries are fairly simple to write and understand, although the language has no visual syntax. Its support of the “natural” graph RDF model, with the simple triple patterns, and a SQL-like syntax, make the language easy to grasp, even for persons non experts of all the intricacies of the semantic web languages and models. Due to the language simplicity, no modularization or extension mechanisms exist for writing complex queries, but the language can be used inside the Java programming language, even mixing it, at a certain extent, with low-level calls to the model’s API.

Functionality—Query Types The language supports only selection and extraction, since the result of a query is a set of bindings based on triple patterns matching and filtering: no kind of “data restructuring” is possible, nor the building of new data. Only the basic data level of RDF is supported, and the programmer must cope with the specialized constructs of the framework like containers, reification, optional properties, as they were ordinary properties. For instance, the following query extracts all the elements from the bag (a kind of container) identified by <http://somewhere.org/bag1>.

```
SELECT ?y
WHERE (<http://somewhere.org/bag1>, ?x, ?y)
AND ! ( ?x eq rdf:type && ?y eq rdf:Bag)
```

The filter part is necessary to eliminate from the result the triple which simply states that the resource is a bag (i.e. has the property `rdf:type` equal to `rdf:Bag`). Queries cannot be compound, and have a single input and output. The filtering part allows the use of simple types URIs, strings, numbers, booleans and the null value, with the corresponding operators. The negation can be used (like in the previous example). On the other hand, the list of triples are only positively conjuncted: no disjunction, negation or optional matching is allowed. This, while severely limiting the expressive power of the languages, has the consequence that a query result is always a set of bindings of values to variables (and not, for instance, subgraphs). Another important limitation is that, although a variable can be bound to a blank node, there is no way to specify in a triple that a node is a blank one, neither with a literal nor with a variable. So, for instance, it is not possible to ask a query which returns all the blank nodes of a graph. No form of recursion or iteration is allowed: only paths of definite lengths can be queried by listing explicitly all the triples forming the path. Finally, no modification to the data can be carried through the language.

Considering the sample queries from Section 3.2, only the first two queries can be expressed in RDQL. Query 1 could be formulated (exactly as in SquishQL above) as

```
SELECT ?essay, ?author, ?authorName
FROM   http://example.org/books
WHERE  (?essay, <books:author>, ?author),
       (?author, <books:authorName>, ?authorName)
USING  books FOR http://example.org/books#,
       rdf FOR http://www.w3.org/1999/02/22-rdf-syntax-ns#
```

Note that everything with an author is considered to be an essay. Otherwise we should add a triple pattern like the following: `(?essay, rdf:type, books:Essay)` if each book is classified as an element of `books:Essay`. Since RDQL is not ontology-aware and no recursion or other mechanism for computing the transitive closure of the subsumption hierarchy is available, it is not possible to select all resources classified as an element of some class that is a sub-class of `books:Essay`.

For Query 2 a similar problem arises. A first version in RDQL could be

```
SELECT ?property, ?propertyValue
FROM   http://example.org/books
WHERE  (?essay, <books:book-title>, "Bellum Civile")
       (?essay, ?property, ?propertyValue),
USING  books FOR http://example.org/books#
```

Note that a property value could be a node with other properties: however, since no recursive mechanism is available in the language, we cannot express a transitive closure of all such properties.

As part of the RDQLPlus (<http://rdqlplus.sourceforge.net/>) implementation of RDQL, an language extension called RIDIQL [259] is defined providing both updates and transparent use of the inference abilities of the underlying Jena Toolkit [126].

Semantics No formal semantics has been published for RDQL.

Complexity and implementation The RDQL has several implementations, of which a well known one is that found in the comprehensive Jena package for semantic web developed in Java at HP Labs [126]. It can work both with an in-memory representation of an RDF model, as well as, for high scalability, with a database based one (currently for MySQL, Oracle E, Postgres). The database representation allows for efficient retrieval of triples by storing them in denormalized tables. The table fields are indexed so that the pattern matching engine can retrieve the triples by using the constants as keys for the search. On the other hand, the filtering part of the query is evaluated in memory on the resulting tuples. Such an approach, while not maximizing the performances of the system, allows the query engine to be implemented with limited complexity.

No formal complexity study of the language has been published so far.

Reasoning No reasoning mechanism is present in the language.

Ontology awareness As already specified, the language ignores every kind of ontological aspect, including typing mechanisms. If such aspects must be considered, they must be treated like all the user-defined data.

Project page:

<http://www.hp1.hp.com/semweb/rdq1.htm>

Implementation(s):

Jena Toolkit [126], RAP (RDF API for PHP) [202], PHP XML Classes (<http://phpxmlclasses.sourceforge.net/>), RDFStore [222], Rasqal RDF Query Library (<http://www.redland.opensource.ac.uk/rasqal/>), Sesame (<http://www.openrdf.org/index.jsp>), 3store (<http://sourceforge.net/projects/threestore/>, cf. [133]), RDQLPlus (<http://rdq1plus.sourceforge.net/>)

Online demonstration:

using Sesame: <http://www.openrdf.org/demo.jsp>

using RAP: http://www3.wiwiss.fu-berlin.de/rdfapi-php/test/custom_rdq1_test.php

using RDFStore: <http://demo.aseantics.com/rdfstore/www2003/>

4.3.1.4 BRQL

BRQL [220] has been recently developed by members of the W3C “RDF Data Access” Working Group as an extension of RDQL [232] aligned with the requirements and use cases detailed in [81]. It is still a very early draft and constantly being improved, therefore this evaluation can only give an impression of the current status.

Several features missing from RDQL but identified as interesting or necessary for an RDF query language in [81] are added, most notably:

- The ability to construct a (single) new RDF graph using the CONSTRUCT keyword. The new graph can be specified with RDQL triple or graph “patterns”.
- A query using the DESCRIBE clause returns the “description” of resources matched by the query part of the expression. The exact meaning of “description” is not yet defined.

- In contrast to RDQL, BRQL supports convenient querying for “quads”, i.e., triples with context information such as source attribution.
- BRQL provides the keyword OPTIONAL to specify triple or graph “patterns” that should be attempted to match but where failure to match does not cause a query solution to be rejected.
- Finally, also a means for testing the non-existence of tuples is added to the language.

For illustrating the capabilities of BRQL, consider again Query 1. But additionally we also would like to return any translator of book, if there is any. This can be expressed in BRQL as

```
SELECT  ?essay, ?author, ?authorName, ?translator
FROM    http://example.org/books
WHERE   (?essay books:author ?author),
        (?author books:authorName ?authorName)
OPTIONAL (?essay books:translator ?translator)
USING   books FOR http://example.org/books#,
        rdf FOR http://www.w3.org/1999/02/22-rdf-syntax-ns#
```

Thanks to the addition of the CONSTRUCT clause, also restructuring and non-recursive inference queries can be expressed. Query 4 can be implemented by the following expression

```
CONSTRUCT (?y books:authored ?x)
FROM      http://example.org/books
WHERE     (?x books:author ?y)
USING    books FOR http://example.org/books#
```

and Query 9 by

```
CONSTRUCT (?x books:co-author ?y)
FROM      http://example.org/books
WHERE     (?book books:author ?x)
          (?book books:author ?y)
AND       (?x neq ?y)
USING    books FOR http://example.org/books#,
```

Project page:

<http://www.w3.org/2004/07/08-BRQL/>

Implementation(s):

none

Online demonstration:

none

4.3.1.5 TriQL

TriQL [40] is under development at the Freie Universität Berlin, German, and aims to extend RDQL to query named graph as introduced in TriG [39] by the authors of TriQL. The reasoning for introducing named graphs in the RDF data model is given in [65].

The use of named graphs allows, e.g., the grouping of assertions by source or author. Then a query such as “Get all books with rating above a threshold of 5. Use only information, that has been asserted by Marcus Tullius Cicero.” can be formulated as

```

SELECT ?books
WHERE ?graph ( ?books books:rating ?rating )
          (?graph swp:assertedBy ?warrant)
          (?warrant swp:authority <http://people.net/cicero>)
USING books FOR http://example.org/books#,
      swp FOR <http://www.w3.org/2004/03/trix/swp-1/>

```

Project page:

<http://www.wiwiss.fu-berlin.de/suhl/bizer/TriQL/>

Implementation(s):

none

Online demonstration:

none

4.3.2 Query Languages influenced by XPath, XSLT or XQuery

In this section, a number of query languages are discussed that have been influenced or are extensions of existing XML query languages developed by the W3C. Some of these approaches (viz. [226, 245, 257]) can be implemented directly on top of these languages merely defining some extension functions and data normalization to be applied before querying. The others propose query languages for RDF that are in spirit and syntax similar to the XML query languages mentioned above.

4.3.2.1 XQuery for RDF: “The Syntactic Web” Approach

During the initial development on XQuery 1.0 [41], Jonathan Robie et al. proposed in a series of articles [224, 226] the use of XQuery for processing RDF. The issue of normalizing RDF before querying is discussed in detail (cf. Section A.3.4). Based on a suitable normal form (essentially statements are considered as triples but statements with same subject are grouped), it is shown how XQuery can be used to query such normalized RDF.

Support for the special semantics of the properties defined in RDFS is added by means of several functions. E.g., the function `rdf:instance-of-class` computes the sequence of all resources (represented by their `description` element) that are an instance of a given class or any of its sub-classes. This is achieved by the following recursive function definition (the first parameter is the set of resources on which to operate):

```

define function
  rdf:instance-of-class($t as element(description)*,
                       $base-name as xs:string)
  as element(description)*
{
  $t[rdf:type = $base-name]
  for $i in $t[rdfs:subClassOf = $base-name]
  return rdf:instance-of-class($t, string($i/@rdf:about))
}

```

Using this function, Query 1 could be formulated (assuming an appropriate normalization has been applied to the RDF data) as:

```
let $t := document("http://example.org/books")//description
for $essay in rdf:instance-of-class($t, "books:Essay"),
  $author in $t[rdf:about = $essay/books:author]
return
  <result>
    {$essay, $author}
  </result>
```

The result of such a query is a sequence of `result` elements containing an essay and one of its author. The name of the author does not need to be queried specifically, since that information is already provided as part of the description of the author selected in the `$author` variable. The query also illustrates that the approach of implementing RDF querying on top of an XML query language has the virtue of being able to return the result of an RDF query in an arbitrary XML format.

The approach also covers the normalization and querying of Topic Maps. Similar to RDF specialized functions are defined to support the specificities of the Topic Maps data model, e.g., the following function computes all derived classes defined in a Topic Maps:

```
define function tm:get-derived-classes($topics as element(topics)*,
  $derivations as element(associations)*, $base as element(topic))
  as element(topics)*
{
  let $a := tm:get-association-by-topic-role($topics,
    $derivations, $base, "superclass")
  for $subclass in tm:get-topic-playing-role($topics,
    $a, "subclass", ())
  return (
    $subclass,
    tm:get-derived-classes($topics, $derivations, $subclass)
  )
}
```

Project page:

none

Implementation(s):

can use any XQuery implementation, however function library has not been made available, some functions are given in [226].

Online demonstration:

none

In [211] it is strongly argued that a unified model of RDF and XML data and a query language based upon such a model is essential for the success of the Semantic Web vision. It is demonstrated, that although the RDF and XML data models differ in some points, a common model theoretic interpretation of RDF and XML data is possible. However, no query language has been proposed based upon this work.

Recently, some information about another approach for extending XQuery for RDF querying, called REX “RDF Extensions to XQuery”, has been discussed in the W3C Data Access Working Group (cf. [234]). This approach seems to be similar to the one discussed above, however RDF statements are generated on-demand by a specific function `related(subject, predicate, object)`. Also no support for RDFS seems to be provided yet. Due to lack of information this extension is not considered further in this survey.

4.3.2.2 XsRQL: An XQuery-style RDF Query Language

XsRQL (XQuery-style RDF Query Language) [149] is a very recent proposal for a RDF query language that borrows from XQuery 1.0 [41], both with respect to the syntax and the design approach. The main objectives of the language design are simplicity and flexibility. In particular, the language aims at providing a syntax flexible enough to allow both the writing of rather simple, concise and more complex but also more expressive queries.

At the core of the proposal are two main differences from XQuery:

- The data model should be adapted to the specificities of RDF. The current draft is rather vague on this point. Some issues can be inferred from the examples given.
- The path language used for accessing and selecting nodes in the data structure has been adapted to the RDF data model: Essentially the same syntax as for XPath is used, however only the `child` axis is supported. Properties are separated from subjects and objects by using the attribute indicator `@` from XPath. However, in contrast to XML attributes, the values (i.e., objects of statements) of RDF properties are not simple, but rather structured values. Therefore after an property further steps may follow in a path expression.

Consider once again Query 1. In XsRQL that query could be stated as

```
declare prefix books: = <http://example.org/books#>;
declare prefix rdf:   = <http://www.w3.org/1999/02/22-rdf-syntax-ns#>;
for $essay in datasource( <http://example.org/books> )/*[@rdf:type/books:Essay],
    $author in $essay/@books:author/*
return
    $essay, $author, $author/@books:authorName/*
```

As XsRQL currently neither supports a closure operation, a descendant-like operator or some other means of traversing an arbitrary-length path in the data structure, it is not possible to return also resources classified by any sub-class of *books:Essay*.

Project page:

<http://www.fatdog.com/xsrql.html>

Implementation(s):

none

Online demonstration:

none

4.3.2.3 XSLT for RDF: TreeHugger and RDF Twig

Similar in spirit to the approaches discussed in Section 4.3.2.1, **TreeHugger** [245] allows the querying and transformation of RDF data in XSLT. However, in contrast to [226] (and due to limitations of XSTL 1.0), the normalization is performed by means of XSLT extension functions and not by an XSLT program. Also the normal form of RDF used for querying is based on the RDF striped syntax [53], but properties are represented both as XML elements and as attributes (raising some problems for multi-valued properties). Three extension functions are provided, one for loading an mere RDF document, one for loading an RDF document and handling the special vocabulary defined by RDFS, and one for loading an RDF document and handling the vocabulary of both RDFS and OWL.

For accessing nodes in an RDF document XPath is used with a special prefix `inv` that allows querying the inverse of a property.

Query 1 could be expressed by the following XSLT stylesheet with TreeHugger extensions:

```
<results xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
  xmlns:books="http://example.org/books#"
  xmlns:th="http://rootdev.net/net.rootdev.treehugger.TreeHugger"
  xmlns:rdfs="http://www.w3.org/2000/01/rdf-schema#"
  xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
  xsl:version="1.0">
  <!-- Load RDF document -->
  <xsl:variable name="doc"
    select="th:documentRDFS('http://example.org/books')" />
  <xsl:for-each select="$doc/books:Essay">
    <xsl:for-each select="books:author/*">
      <result>
        <xsl:value-of select="inv:books:author" />
        <xsl:value-of select="." />
        <authorName>
          <xsl:value-of select="books:authorName/*" />
        </authorName>
      </result>
    </xsl:for-each>
  </xsl:for-each>
</results>
```

Project page:

<http://rdfweb.org/people/damian/treehugger/>

Implementation(s):

available from the project page

Online demonstration:

<http://swordfish.rdfweb.org/discovery/2003/09/treehugger/>

In [257], another approach of extending XSLT 1.0 with functions for querying RDF, called **RDF Twig**, is described. In contrast to the previously discussed proposals, it provides different views on the RDF data corresponding to redundant or non-redundant (i.e., where nodes that are reachable by various paths are repeated, resp. not repeated) depth or breadth first traversals of the RDF graph. Furthermore, two query mechanisms are provided: A small set of logical operations on the RDF graph (also used in the example below) and an interface to the RDQL query engine provided by the Jena Toolkit [126] used for implementing RDF Twig.

To give a feeling for the language, we consider once more Query 1 and how it can be realized in this query language:

```
<xsl:stylesheet xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
  version="1.0"
  xmlns:rt="http://nwalsh.com/xslt/ext/com.nwalsh.xslt.saxon.RDFTwig"
  xmlns:twig="http://nwalsh.com/xmlns/rdftwig#"
  xmlns:books="http://example.org/books#"
  xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#">
  <xsl:template match="/">
    <xsl:variable name="model"
      select="rt:load('http://example.org/books')"/>
    <!-- this is used as default model from now on-->
    <xsl:variable name="pType"
      select="rt:property('http://www.w3.org/1999/02/22-rdf-syntax-ns#','type')"/>
    <xsl:variable name="essays"
```

```

        select="rt:find($label, 'books:Essay')"/>
<xsl:variable name="tree"
              select="rt:twig($essays)/twig:result"/>
<results>
  <xsl:for-each select="rt:find($label, 'books:Essay')">
    <result>
      <xsl:value-of select="rt:twig(.)" />
      <xsl:value-of select="rt:twig(.)twig:result/books:author" />
    </result>
  </xsl:for-each>
</results>
</xsl:template>

```

For simplicity, only essays and authors are considered, the names of the authors will be returned as well, as they are reachable from the associated essay and `rt:twig(.)` returns all information reachable from the current essay (in this case). RDF Twig does not support RDFS or OWL, therefore only resources classified directly as `books:Essay` will be considered by this query.

Project page:

<http://rdftwig.sourceforge.net/>

Implementation(s):

available from the project page

Online demonstration:

none

4.3.2.4 RDFT and Nexus Query Language: XSLT-style RDF Query Languages

RDFT As XsRQL is modeled after XQuery, RDFT [88] is a draft proposal closely related to XSLT 1.0. As XSLT 1.0 it uses templates that are matched recursively against the data structure. Naturally, the structural recursion is performed against an RDF graph, raising issues with cyclic graph structure that are still open issues in the development of RDFT.

RDFT uses an adaptation of XPath for querying RDF graphs, called NodePath. As in most of the other approaches oriented on XML query languages, a striped view of RDF [53] is adapted where properties and other resources alternate. No provision for querying data described with RDFS or OWL is made.

Only a subset of XSLT elements is supported, but a macro mechanism is introduced as illustrated in the following implementation of Query 1 in RDFT:

```

<rt:stylesheet rt:version="1.0" xmlns:rt="http://purl.org/vocab/2003/rdft/">
  <rt:macro-set rt:prefix="rdf">
    <rt:macro name="type"
              value="resource('http://www.w3.org/1999/02/22-rdf-syntax-ns#type')/resource()"/>
  </rt:macro-set>
  <rt:root-template>
    <rt:apply-templates
      rt:select="/resource()[rdf:type = resource('http://example.org/books#Essay')]/>
  </rt:root-template>
  <!-- Template for the Essay
  <rt:template pattern="resource()[rdf:type =
    resource('http://example.org/books#Essay')" />
    <xsl:value-of select="." />
  <rt:apply-templates

```

```

      rt:select="resource('http://example.org/books#author')/resource()"/>
    </rt:template>
    <!-- Template for the author -->
    <rt:template
      pattern="resource('http://example.org/books#author')/resource()">
      <xsl:value-of select="." />
    </rt:template>
  </rdf:stylesheet>

```

Again, for simplicity, only books and their authors are returned without considering the names of the authors. Also note, that the specification is not really clear what the result of such a query will be: an XML tree or some form of an RDF graph. The description of `rt:element` seems to indicate the former, the description of `rt:value-of` the latter.

Project page:

<http://www.semanticplanet.com/2003/08/rdft/spec>

Implementation(s):

none

Online demonstration:

none

Nexus Query Language In [85] another approach for querying RDF (and some form of XML) using an XSLT-like language is sketched. The basic idea is to translate both RDF/XML and also some non-RDF XML documents into a hierarchy of elements (that also carry some attributes) based upon the relations between the elements. The result of a query is then some (also hierarchical) view over this element tree. [85] gives no consideration w.r.t. cyclic relations among elements but the language used seems to indicate that only proper hierarchies can be represented.

RDF statements are mapped to nodes in the data model in the following way: nodes in the graph represent RDF properties; an RDF statement (S, P, O) is represented by edges from all nodes representing some property with the value S to a node representing the property P with value O . A resource that never occurs as object is assigned as value to a special property called `query:seed`. [85] seems to indicate that there can be only one such `query:seed` node, an assumption that is clearly invalid for general RDF graphs.

The query language provides means for matching such property nodes based on the identifier (represented as URI or XML QName) of the property and the type (as determined by an `rdf:type` statement) of the value of the property.

Consider again Query 1 and the following Nexus query that implements the query:

```

<query:plan>
  <query:template match="query:seed" type="books:Essay">
    <query:call name="query:insert" rename="book">
      <query:call name="query:format" rename="title"
        value="book:title" />
      <query:call name="query:traverse" />
    </query:call>
  </query:template>
  <query:template match="book:author">
    <query:call name="query:insert" rename="author">
      <query:call name="query:format" rename="name"
        value="book:authorName" />
    </query:call>
  </query:template>
</query:plan>

```

```
</query:call>
</query:template>
</query:plan>
```

An excerpt of the result of this query on the sample data from Figure 2.1 would be:

```
<book title="Bellum Civile">
  <author name="Julius Caesar" />
  <author name="Aulus Hirtius" />
</book>
```

Obviously, the syntax is rather verbose and does not inherit the ability of XSLT to write arbitrary XML as content of an element. Furthermore, the development of the language seems to be stalled as the only information available on this language is the rather short and often vague report cited above [85].

Project page:

none

Implementation(s):

not publicly available, no report on any implementation

Online demonstration:

none

4.3.2.5 XPath-style Access to RDF: RDF Path, RPath and RxPath

Several mostly sketchy proposals [247, 89, 208, 185] for adapting the XPath-style navigational access to RDF graphs have been published in recent years. For this survey, two representative ones have been selected, viz. “Pondering RDF Path” by Sean Palmer and RPath [185], a language designed in the context of device independence and content adaptation.

RDF Path In [208], a sketch of an RDF Path language is proposed that is closely aligned with XPath. The aim of this language is to provide clear equivalents for the XPath facilities such as selection of context nodes, filtering and location steps, but operating on RDF data. The syntax is similar to XPath but extended by special node-tests for RDF, such as `arc()` and `subj()` for selecting all arcs, resp. all subjects in the RDF data. Several aspects of XPath not relevant for RDF are dropped. Only basic navigational features are provided, functions and testing of values is mostly not considered in this early draft. Finally, as in most of the approaches based on XML query languages, the fact that, in contrast to XML trees, RDF graphs are not rooted is not considered.

As XPath, this language is not capable to select related information as in Query 1. Therefore a slightly variation of Query 2 is used to illustrates its abilities: “Select the names of all authors of historical essays with the title ‘Bellum Civile’”. This query can be realized by the path expression

```
*[rdf:type/books:Historical_Essay books:title/"Bellum Civile"]/
books:author/*/books:authorName
```

Appropriate mappings for the prefixes used in the query have to be established before evaluating the query (this is also the case in XPath). Also note that due to the lack of support for the special semantics of the vocabulary provided by RDFS, this query will only return resources directly classified as `books:Historical_Essay`.

Project page:

<http://infomesh.net/2003/rdfpath/>

Implementation(s):

none

Online demonstration:

none

RPath [185] is a new RDF query language which is based on the *path navigation* principle known from XPath. In fact, the major design idea behind RPath is to provide *XPath for RDF*. In its current state, however, it focuses strongly on CC/PP and UAProf, two RDF applications for describing device characteristics like *color capable*, *color depth*, or *screen resolution*. CC/PP is the general framework for those *device profiles*, whereas UAProf is a specific vocabulary focused on (but not limited to) mobile devices.

As RPath navigates through RDF data using paths, it views this data as graph, not as triples. The concepts of the language resemble very much that of XPath, that is, location steps, *vertex-edge-tests* (corresponding to node-tests in XPath), and predicates. Differences to XPath are due to the differences between the data models of XML and RDF, for example: The axes can follow a path along vertices (RDF predicates) and edges (RDF subjects and objects). The adaptation of most XPath concepts to the RDF data model is straightforward. One major difference, however, is the absence of a root node in the RDF graph, and the question of finding a start point for the path expression remains an open topic.

On the other hand, the (current) focus of RPath lies on CC/PP and UAProf, which limit RDF graphs to rooted two-level trees. The language itself does not offer any CC/PP-specific or UAProf-specific features, all of it can be used to query generic RDF graphs. The implementation of the prototype, however, is capable of handling the protocol specified by CC/PP, including *default profiles* and *profile diffs*, as well as the data types defined by UAProf.

The authors of RPath claim its ease of use as the main advantage. As XPath already enjoys widespread use in the XML world, the learning curve for RPath should be shallower than that for other RDF query languages for potential users. Another advantage is the tight coupling with CC/PP and UAProf, which should make RPath suitable to be used in *device independence*-applications.

RPath has been developed at Keio University, Japan, by Keita Matsuyama, Michael Kraus, Kazuhiro Kitagawa (Activity lead for the W3C Device Independence), and Nobuo Saito. Currently, development has been halted.

To convey an impression of the languages capabilities and syntax, the same variation of Query 2 as above is used: "Select the names of all authors of historical essays with the title 'Bellum Civile'."

```
/@vertex()[  
  rdf:type/@books:Historical_Essay and
```



```
books:title/@vertex()[equals('Bellum Civile')]
]/books:author/books:authorName
```

Note, that in contrast to most other path-based approaches to querying RDF data, RPath does not require the user to write paths where expressions matching vertices (i.e., classes) and edges (i.e., properties) alternate (similar as in striped RDF [53]). This is possible, as all steps begin with an axis specification and different axes for vertices and edges are provided. To illustrate this point the same query shown above could be written using non-abbreviated RPath syntax as:

```
outerVertex::vertex()[
  outEdge::rdf:type/outVertex::books:Historical_Essay and
  outEdge::books:title/outVertex::vertex()[equals('Bellum Civile')]
]/outEdge::books:author/outEdge::books:authorName
```

Project page:

none

Implementation(s):

prototype in Java, based on a CC/PP engine from Sun

Online demonstration:

none

RXPath As part of the Rx4RDF project <http://rx4rdf.liminalzone.org/rx4rdf> that aims at improving the accessibility of RDF for non-experts another adaption of XPath for querying RDF data has been defined.

In contrast to the approaches discussed above and somewhat related to TreeHugger and RDF Twig, RxPath is essentially “ a mapping between the RDF Abstract Syntax to the XPath Data Model” [241]. The mapping consists in four steps:

- One top-level element in the XML document is created for every resource in the RDF model with the type of the resource as element label.
- “Each root element has a child element for each statement the resource is the subject of. The name of each child is [the] name of the property in the statement.” [240]
- “Each of these children have [a] child text node if the object of the statement is a literal or a child element if the object is a resource.” [240]
- “Object elements have the same name and children as the equivalent root element for the resource, thus defining a potentially infinitely recursive tree.” [240]

As stated, such a mapping might lead to infinite trees, in particular when evaluating any of the closure axes of XPath (descendant, following, preceding, etc.) the number of nodes selected in the tree is no longer finite. RxPath proposes a circularity-test for the evaluation of such axes, such that whenever an element with the same URI reference as an ancestor is encountered that element is skipped in the evaluation. (One consequence of this approach is that blank nodes need to be assigned a unique URI reference.)

Furthermore, RxPath changes the semantics of the closure axes to only consider elements representing RDF properties in the original RDF model (this is easy as the mapping from RDF

into an XML document discussed above uses a striped representation of RDF statements [53]). Furthermore, an expression such as `descendant::rdf:type` only matches an element representing an `rdf:type` property where all elements representing an RDF property actually represent an `rdf:type` property. In other words, `descendant::rdf:type` is more similar to the regular tree expression `(rdf:type)*` than to the XPath expression `descendant::rdf:type`.

Once more we use the same variation of Query 2 as above to illustrate the language syntax: “Select the names of all authors of historical essays with the title ‘Bellum Civile’.” (assuming the books prefix is bound to `http://example.org/books-rdfs#`):

```
/books:Historical_Essay[books:title = 'Bellum Civile']/  
  books:author/*/books:authorName
```

Based on RxPath two more languages have been defined: RxSLT [242] is “syntactically identical to XSLT 1.0” [242], but uses RxPath instead of XPath 1.0. RxUpdate [243] is syntactically very similar to XUpdate [168], but again uses RxPath instead of XPath to update RDF models.

Project page:

<http://rx4rdf.liminalzone.org/rx4rdf>

Implementation(s):

prototype in Python, available from project page

Online demonstration:

none

4.3.2.6 Versa

Conceived as query language for the Python-based 4Suite² toolkit for XML and RDF application development, Versa³ is a query language for RDF inspired by XPath that can be used as a replacement of XPath for pattern matching in XSLT. Although inspired by XPath it is sufficiently different to deserve a discussion separate of the languages shown in the previous section.

The details of Versa are described in [204], [201] and [200] present gentle introductions into the language. The core design principles of Versa, as stated on <http://uche.ogbuji.net/tech/rdf/versa/>, are:

- **“Strong alignment with XML.”** As the 4Suite toolkit provides access to both XML and RDF data and technologies, the use of Versa to access RDF, e.g., when constructing an XML document with XSLT is an obvious choice. However, it has not been attempted to provide a single query language for both XML and RDF data, rather a set of query languages such as XPath and Versa are provided, each specialized for a certain data formalism.
- **“XPath-like idiom.”** The approach taken in the 4Suite toolkit, viz. not to provide an integrated query language for different data formats but rather to use a set of specialized query languages is all the more viable the more these query languages have in common. Therefore, Versa has been designed with a syntax inspired by XPath, although it arguably deviates quite notably, even more so than the query languages discussed in the previous section.

² <http://4suite.org/> ³ <http://uche.ogbuji.net/tech/rdf/versa/>

- **“Extensibility.”** Just as XPath, Versa is designed to be extensible in the same way as XPath, i.e., using externally defined extension functions. However, the current version of the specification [204] is not very clear on this point.
- **“Ease of learning.”** The authors claim that “many users have reported that they become proficient very quickly with Versa.” Justifiably, they argue that the superficial similarity some of the other RDF query languages share with SQL is not actually helpful in many cases as there is considerable mismatch with regard to the data model used by the languages. Although the traversal constructs of Versa are designed to be easy to recognize and remember, they are considerable different to the traversal expressions in XPath (i.e., XPath axes and node-tests) or similar path-based query languages.
- **“Expressiveness.”** The set of query constructs provided by Versa covers a, for a path-based language to be used within, e.g., XSLT, surprisingly large set of the functionalities discussed in Section 3.2, lacking most notably means for defining views, functions, and other forms of construction. This lack can be justified to some extent by the intended use of Versa within some host language that might provide these means.

At the core of the Versa query language are the assorted traversal and filter expressions.

- *Forward traversal.* Versa allows the traversal of one or more properties starting from a list of subjects to select the objects that are reachable via the given properties. E.g., the expression `all() - books:author -> *` selects all resources that are author of another resource. The objects can also be restricted, e.g., to those containing a certain string. Here `*` indicates a wildcard, i.e., no further restriction on the objects. Such traversal expressions can be chained.

The following Versa query uses forward traversal operators to implement Query 1 (in the following discussion of Versa, the namespaces are assumed to be set up externally):

```
distribute(type(books:Essay),
           "distribute(.-books:author->*,
                    ".", "-books:authorName->*)")
```

The `distribute()` function returns a list of lists containing the result of the second, third, ... argument evaluated starting from each of the resources selected by the first argument. As in XPath, `.` denotes the current node in such a context. Here, the first argument selects all resources classified as `books:Essay` and evaluates starting with these resources the remaining two arguments: The former returns all those books, the second uses `distribute()` again to select the authors together with their name.

- *Forward filter.* Just like a traversal, where the object of the traversed statement is select, one can use a forward filter to select the subject of a statement. E.g., the Versa query `type(books:Essay) |- books:title -> eq("Bellum Gallicum")` selects all essays with a title “Bellum Civile”.
- *Backward traversal.* Sometimes, one would like to navigate from the objects of a statement to the subjects. Therefore, Versa offers also a backward traversal (although, so far no backward filter is provided, it can however, be implemented with the general filter expression discussed below). E.g., the query above selecting all essays with a title “Bellum Gallicum” can also be written as

```
(books:Essay <- rdf:type - *) |- books:title -> eq("Bellum Gallicum")
```

- *General traversal.* Whereas the traversal operators discussed so far only allow the traversals of paths with fixed length, Versa also offers a function for general traversals, both forward and backward. This function, called `traverse`, can also be used to traverse paths of arbitrary length. E.g., the following query obtains all sub-classes of `books:Writing`:

```
traverse(books:Writing, rdf:subClassOf, vtrav:inverse, vtrav:transitive)
```

- *General filter.* Similarly, the `filter` function provides a general filter, where the result of evaluating the first argument is filtered by the remaining arguments evaluated in the context of the elements selected by the first one. E.g., to select all essays with title “Bellum Gallicum” and a translator with name “J. M. Carter” could be implemented by the following query:

```
filter(books:Essay <- rdf:type - *,  
      ". - books:title -> eq('Bellum Gallicum')",  
      ". - books:translator -> books:translatorName -> eq('J. M. Carter')")
```

Following this short overview over the core traversal and filtering functions of Versa, the language is further investigated following the criteria proposed in Section 3.2.

Easy of use. Although designed to be closely aligned with XPath, the traversal operators shown above and the use of functions instead of specific syntactical constructs such as predicates in XPath gives Versa a rather unfamiliar feeling compared to other XPath-oriented languages.

Functionality—Query Types. Selection and extraction queries can be easily implemented in Versa, although as demonstrated by the above implementation of Query 1, the selection of multiple related items is not very convenient. In contrast to most other RDF query languages, Versa allows the extraction of arbitrary size graphs, as required by Query 2. Also reduction queries can be expressed, e.g., using negation or set difference. Query 3 can be implemented in Versa by the following query

```
difference(all(),  
          union(type(rdfs:Class),  
                union(type(rdfs:Property,  
                        all() <- books:translator - *)))  
          )  
        )
```

This query selects all resources except for those that are either (a) a `rdfs:Class`, (b) a `rdfs:Property`, or (c) occur as object in a statement with predicate `books:translator`.

Neither restructuring nor combination or inference queries can be directly expressed in Versa, as the result of a Versa query is always a list (or possibly a list of lists). However, queries such as Query 4 and 9 can be approximated, e.g., by returning all the tuples of co-authors of a book:

```
distribute(all(), ". - books:author -> *", ". - books:author -> *")
```

However, this query will also include that, e.g., "Julius Caesar" is a co-author of himself. It does not seem possible to avoid this in Versa, as the later arguments of `distribute` are evaluated independently.

Versa also provides a large set of aggregation functions. Query 5 can be implemented by the following program

```
max(filter(all(),
  ". - books:author -> books:authorName -> eq('Julius Caesar')"
),
  - books:year -> *)
```

Starting from the filtered books, the year of publication of those books is selected and the maximum of these years calculated.

Also Query 6 can be implemented in Versa using the `length` function for calculating the number of authors per book:

```
distribute(traverse(books:Writing, rdf:subClassOf,
  vtrav:inverse,vtrav:transitive),
  ". - books:author -> books:authorName -> eq('Julius Caesar')",
  "max(length((. <- rdf:type *) - books:author -> *))"
)
```

Semantics. No formal semantics has been provided so far.

Complexity and implementation. No formal complexity study of the language has been published so far. Only a single implementations by the authors of the language is available.

Reasoning. No reasoning abilities are provided.

Ontology awareness. Versa provides a RDFS-aware `type()` function that returns all resources that are classified under the given RDFS class or one of its sub-classes. The transitive semantics of `rdfs:subClassOf` and `rdfs:subPropertyOf` are not provided by default but can be implemented by the general traversal function `traverse`.

Project page:

<http://uche.ogbuji.net/tech/rdf/versa/>

Implementation(s):

available as part of 4Suite from <http://4suite.org/>

Online demonstration:

none

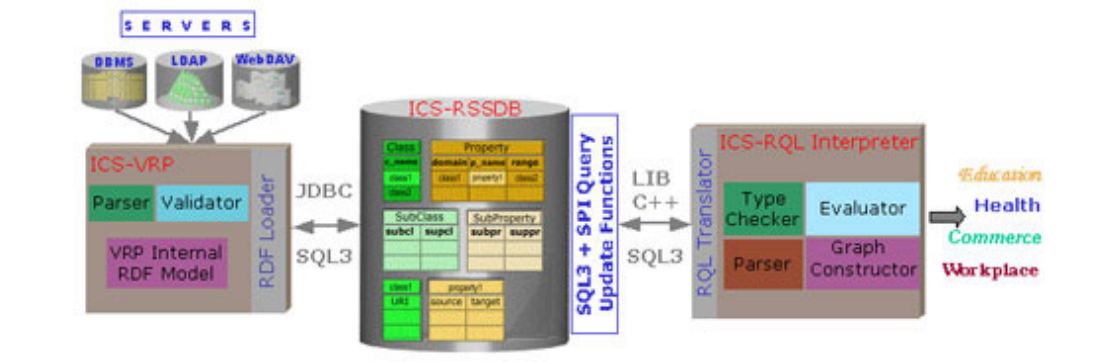
4.3.3 RQL-family

4.3.3.1 RQL

RQL [76, 146, 145, 147, 148], the RDF Query Language, has been developed at the ICS-FORTH and influenced several later proposals for RDF query languages such as SeRQL (cf. Section 4.3.3.2, eRQL (cf. Section 4.3.3.3, and BRQL (cf. Section 4.3.1.4). The reference implementation of RQL has been developed as part of the RDFSuite [9], a collection of tools that provide efficient access to increasingly large RDF stores drawing on established database technology provided by an ORDBMS. The architecture of RDFSuite is shown in Figure 4.1 identifying the three main components of RDFSuite:

- the Validating RDF Parser (VRP), a high-performance, RDFS-aware RDF parser that also allows the user to specify semantic constraints against which a document should be validated,
- the RDF Schema-Specific Data Base (RSSDB), a persistent RDF store based upon the ORDBMS PostgreSQL (<http://www.postgresql.org>) that stores RDF data based on its (RDFS) schema, and
- an interpreter for RQL, implemented on top of the RSSDB.

Figure 4.1 RDFSuite Architecture
(from [8], © ERCIM News)



RQL has also been used in the ICS-FORTH Semantic Web Integration Middleware (SWIM) [75], where it is used to query data represented as RDF but integrated from different data sources, e.g., XML documents or relational databases.

In contrast to RDF query languages such as RDQL, that are tailored to be easy to use by providing only a small set of (often used) query constructs excluding, e.g., the direct exploitation of RDFS or direct support for more complex aspects of RDF such as containers or reification, RQL also has the ability to combine schema and data querying and allows the specification of complex graph patterns.

At the core of RQL is a formal data model for RDF graph data (deviating slightly from standard RDF/S semantics by disallowing cycles in the subsumption hierarchy and requiring that for each property both domain and range are defined for (a) simplification and (b) alignment

with underlying type systems) based upon typing information provided by an RDFS schema. The salient features of this data model are

- a clear separation between the different RDF/S abstraction layers: (1) data, i.e., description of resources such as persons, XML documents, etc., (2) schemas, i.e., classifications for such resources, and (3) metaschemas containing the metaclasses (i.e., classes of classes or properties) `rdfs:Class` and `rdfs:Property` and their refinements)
- flexible type system tailored to the specificities of RDF/S by allowing (a) optional and multi-valued properties, (b) superimposed descriptions of the same resources (i.e., resources can be classified under different classes not related in the subsumption hierarchy) and (c) the flexible refinement of schemas.

Based upon this data model, RQL provide a number of novel query constructs for querying the type information associated with the RDF data. In the following a condensed overview of the most prominent query constructs provided in RQL given, more details can be found in [76, 146, 145, 147, 148].

Basic schema queries. As stated above, one of the salient features of RQL is the use of type information derived from an RDFS schema. The subsumption hierarchy defined by `rdfs:subClassOf` and `rdfs:subPropertyOf` can be accessed in different manners:

- Querying the *sub-classes* of a class: E.g., `subClassOf(books:Writing)` returns all sub-classes of `books:Writing` (assuming the namespace `books` is set up properly using `USING NAMESPACE books = &http://example.org/books#`).
- Querying the *sub-properties* of a property using `subPropertyOf`.
- Querying the **domain and range of a property**. E.g., the following query obtains instances of which classes can be combined with `books:author`:

```
SELECT $C1, $C2
FROM   {$C1}books:author{$C2}
USING  NAMESPACE books = &http://example.org/books#
```

The `$` prefix of a variable indicates a “class variable”, i.e., a variable ranging on schema classes, in other words, resources with `rdf:type rdfs:Class`. Therefore, `$C1` selects all classes that are in the domain of `books:author` and `$C2` those in its range.

More explicitly, the query can be stated as (however, since `subClassOf` is not reflexive, the direct domain and range are left out):

```
SELECT C1, C2
FROM   subClassOf(domain(book:author)){C1}, subClassOf(range(books:author)){C2}
USING  NAMESPACE books = &http://example.org/books#
```

Note, that here the query variables are *not* prefixed by `$`, therefore the values (and not their type) returned by the `subClassOf` expressions are used. This is the intended result, as the values returned by `subClassOf` are already the classes searched for.

Finally, the query could also be formulated using a type constraint:

```
SELECT C1, C2
FROM   Class{C1}, Class{C2}, {;C1}books:author{;C2}
USING  NAMESPACE books = &http://example.org/books#
```

The first part of the FROM clause (ClassC1, ClassC2) selects all pairs of classes defined in the schema. These pairs are further constrained by the second part, that stipulates that only such pairs are to be retained where the first class can occur as domain and the second as range of books:author. This restriction is expressed by a type constraint, written (in general) {X;Y}, where X is bound to the concrete values (here, e.g., “Julius Caesar”) and Y to the type of the value (here, e.g., Literal, i.e., the class of all RDF literals).

- Querying **tops and leafs** of the subsumption hierarchy. E.g., topclass(books:Historical_Essay) returns books:Writing on the data from Figure 2.1. Furthermore, it is possible to query the nearest common ancestor of two classes, e.g., nca(books:Historical_Essay, books:Historical_Novel) is books:Essay.
- Querying **meta-schema** information. Meta-classes, such as Class, and Property (the class of all classes, all RDF literals, all properties, resp.) can be queried just as any other class. Obviously, the extent of a meta-class is a set of classes, as demonstrated above.
- Querying RDF **properties**. Just as classes can be queried using class variables indicated by a \$-prefix, RDF properties are selected by “property variables” prefixed by @. E.g., the following query selects all properties together with their range that can be attached to resources classified as books:Writing:

```
SELECT @P, $V
FROM   {;books:Writing}@P{$V}
USING  NAMESPACE books = &http://example.org/books#
```

Combining these facilities, one can easily implement Query 8 in RQL:

```
SELECT X, Y
FROM   Class{X}, subClassOf(X){Y}
```

This query returns pairs of all classes such that the first class is a super-class of the second one.

Data queries. Obviously, RQL also provides access to the actual resource descriptions. One can access resources by their type and by navigating to their position in the RDF graph, and further restrictions to the data to be selected can be specified by filters:

- Querying the *extent* of classes and properties. Both classes and properties can be queried for their (direct and indirect) extent:

```
books:Writing
```

returns all resources classified as books:Writing or one of its sub-classes. This is equivalent to

```
SELECT X
FROM   books:Writing{X}
```

If only those resources X shall be returned that are directly classified as `books:Writing`, i.e., where an RDF tuple $(X, \text{rdf:type}, \text{books:Writing})$ exists, the name of the class must be prefixed by \wedge .

Similarly, one can query the extent of a property, e.g.,

```
 $\wedge$ books:author
```

returns pairs of all resources X, Y that stand in the `books:author` relation, i.e., where an RDF tuple $(X, \text{books:author}, Y)$ exists.

- Querying by navigation using **generalized path expressions** [74]. RQL uses generalized path expressions known from, e.g., OQL and Lorel, to allow navigation both in the data and in the schema graph.

This allows for an easy traversal of relations convenient for implementing Query 1:

```
SELECT X, Y, Z
FROM   {X;books:Essay}books:author{Y}.books:authorName{Z}
USING  NAMESPACE books = &http://example.org/books#
```

- **Filtering** the result. The result of a query can be further restricted by a WHERE clause. E.g., to select only books and the names of their authors if the title of the book is "Bellum Civile", one could add a WHERE clause as in the following query:

```
SELECT X, Y
FROM   {X;books:Essay}books:author.books:authorName{Y}, {X}books:title{T}
WHERE  T = "Bellum Civile"
USING  NAMESPACE books = &http://example.org/books#
```

Mixing schema and data queries. Querying data and schema can be intertwined in any way. In particular, as shown above, data can be selected based on its type, by providing filter expressions for the variable selecting the data. E.g., the expression `X;books:Essay` restricts bindings for variable X to resources with type `books:Essay`.

Often interesting queries not only require the use of type information for filtering, but also benefit from providing type information in the result. Query 2, that asks for a kind of description of the book with title "Bellum Civile", could be implemented in RQL in the following way (interpreting "description" as the schema of a resource):

```
SELECT $C, (SELECT @P, Y
            FROM   {Z ;  $\wedge$ $D}  $\wedge$ @P {Y}
            WHERE  Z = X and $D = $C)
FROM    $\wedge$ $C {X}, {X}books:title{T}
WHERE  T = "Bellum Civile"
USING  NAMESPACE books = &http://example.org/books#
```

This query returns all classes under which the resource with title "Bellum Civile" is directly classified (\wedge `$C{X}` selects all values in the direct extent of any class). Grouped by these classes,

all properties and classes that can be used as values for the respective property are queried by a nested RQL query and returned.

Several of the features of RQL are not discussed in length here, such as support for containers, aggregation, schema discovery, etc.

Although the original proposal of RQL does not include view definition constructs, RVL [175] provides such an extension. Using this language, one could, e.g., define the inverse relation of `books:author` as a view by the following program:

```
CREATE NAMESPACE mybooks = &http://example.org/books-rdfs-extension#
VIEW   authored(Y, X)
FROM   {X}books:author{Y}
USING  NAMESPACE books = &http://example.org/books#
```

Such a view can be queried just like any other data, e.g., to select all authors that authored a book translated by someone called “Carter” together with these books:

```
SELECT X, Y
FROM   {X}mybooks:authored{Y}books:translator{T}
WHERE  T like "*Carter*"
USING  NAMESPACE books = &http://example.org/books#
USING  NAMESPACE mybooks = &http://example.org/books-rdfs-extension#
```

Concluding this overview of RQL is a short discussion of the query language under the evaluation criteria proposed in Chapter 3.

Easy of use. The comparatively large and diverse set of features offered by RQL, naturally lead to a somewhat more complicated syntax and semantics if compared to basic query languages such as RDQL (cf. Section 4.3.1.3). One critique voiced when considering RQL is that the typing features of the language, although useful for some scenarios, actually complicate the expression of simple queries and are hard to implement. Therefore, both of the proposals for simplifying RQL, viz. SeRQL and eRQL, do not provide the same level of type information as RQL.

Also the use (and representation) of a number of syntactic constructs, e.g., for distinguishing class (using `$`-prefix), property (using `@`-prefix), and data variables (using no prefix) or for separating direct (`^`-prefix) and full extent (no prefix) of a class, is neither intuitive nor based on established conventions. Providing also a more verbose version of such constructs, might improve readability and ease the familiarization with the language.

Functionality—Query Types. As stated already, RQL is far more expressive than basic RDF query languages such as the SquishQL family. Actually, most of the queries discussed in Section 3.2 can be easily expressed in RQL (here, we also include the view definitions provided by RVL) with the notable exception of those queries requiring means for traversing the transitive closure of arbitrary relations, instead of only the two relations `rdfs:subClassOf` and `rdfs:subPropertyOf` that constitute the subsumption hierarchy.

An implementation of Query 1 has already been given above. Query 2 can not be expressed in RQL exactly, as there is no means to select “everything related to some resource”. However, in the above discussion a modified version of this query, where a resource is described by its schema, is shown. Reduction queries such as Query 3 can often concisely be expressed in RQL, in particular, if the reduction query is, as in this case, based on type information:

```

SELECT S, @P, 0
FROM   (Resources minus (SELECT T FROM {B}books:translator{T})) {S},
       (Resources minus (SELECT T FROM {B}books:translator{T})) {S},
       {S}@P{0}
USING  NAMESPACE books = &http://example.org/books#

```

This query returns all triples where both the subject and the predicate of the triple is some Resource that does not occur as object in a statement with predicate `books:translator`. Note, that Resource is the top-level *schema* class, but does not include resources classified as `rdfs:Class` or `rdfs:Property` as these are considered schema objects and therefore contained in a meta-schema class.

An implementation of Query 4 is given above using RVL.

Aggregation queries are also supported by RQL, e.g., consider Query 5 and a possible implementation in RQL shown below:

```

max(SELECT Y
FROM   {B;books:Writing}books:author.books:authorName{A},
       {B}books:pubYear{Y}
WHERE  A = "Julius Caesar")

```

Inference and combination queries often require traversing the transitive closure of arbitrary relations or even general recursion, that can both not be expressed in RQL. However, when the inference does not need recursion, as in Query 9, the query can be expressed in RQL

```

SELECT A1, A2
FROM   {Z}books:author{A1}, {Z}books:author{A2}
WHERE  A1 != A2
USING  NAMESPACE books = &http://example.org/books#

```

or in RVL (allowing the results to be queried as any relation in the original data)

```

CREATE NAMESPACE mybooks = &http://example.org/books-rdfs-extension#
VIEW   mybooks:co-author(A1, A2)
FROM   {Z}books:author{A1}, {Z}books:author{A2}
WHERE  A1 != A2
USING  NAMESPACE books = &http://example.org/books#

```

Semantics. RQL has been designed around a formal data model for RDF. Based upon this data model, both typing rules and formal interpretations of RQL queries have been specified [148]. As discussed above, the data model slightly deviates from a standard RDF interpretation (as specified in [167]).

Complexity and implementation. No formal complexity study of the language has been published so far. Aside of the implementation provided by the authors of the language as part of the ICS-FORTH RDFSuite, the only other implementation of RQL the authors of this survey are aware of is as part of Sesame [59]. However, for Sesame the typing features of RQL have been disregarded for the most part and a new RDF query language, called SeRQL, that is based upon RQL has been specified. It is discussed in the following section.

Reasoning. No specific reasoning abilities are provided. In particular, the only form of recursion is the ability to navigate the transitive closure of certain predefined schema relations. However, some non-recursive inference queries can be realized as shown above. In particular, RQL supports a rich set of boolean expressions including negation and quantification.

Ontology awareness. RQL strongly relies on RDFS for providing type information on the data to be queried. (An extension to) OWL has not been considered so far.

Project page:

<http://139.91.183.30:9090/RDF/RQL/>

Implementation(s):

RDFSuite (<http://139.91.183.30:9090/RDF/index.html>), Sesame (<http://www.openrdf.org/>) without type system

Online demonstration:

<http://139.91.183.30:8999/RQLdemo/>, based on Sesame: <http://www.openrdf.org/demo.jsp>

4.3.3.2 SeRQL

As part of the EU research project On-To-Knowledge (<http://www.ontoknowledge.org/>), the Sesame [59] RDF database has been developed. An initial review of query languages for semi-structured data and RDF [56] suggested the use of RQL as query language for the On-To-Knowledge project. In [57], this choice is detailed by defining the query language to be used in the On-To-Knowledge project, tentatively called OTK-RQL. This language is still syntactically nearly identical to RQL, but leaves out in particular RQL's entire type system. The current version of Sesame still supports RQL (without typing of queries), but also provides a novel RDF/RDFS query language claiming to “combine the best features of other (query) languages (RQL, RDQL, N-Triples, N3)” [63] into a “second generation RDF query language” [58].

SeRQL is described in detail in [63] and [58]. Here, we focus on the differences to previous RDF query languages, in particular to RQL.

The most striking differences between RQL and SeRQL are:

- SeRQL does not provide any form of **typing** for RDF resources, only typed literals as provided by the recent revision of RDF [197] are considered. One reasons the others of SeRQL offer for avoiding to provide typing in the query language (aside of the increase complexity of both implementation and language) is that RDF schema languages are designed to be extensible and a typed query language therefore needs to provide some mechanism for integrating such extensions (as provided, e.g., by OWL), making an implementation of such a query language even harder.
- In an effort to simplify common queries, SeRQL modifies and extends the general **path expressions** used in RQL: Basic path expressions use a syntax similar to RQL, e.g., `{X} <rdf:type> {<books:Essay>}` is the SeRQL notation for a path expression that returns all resources classified under `books:Essay` as bindings for the variable X. Compound path expressions, e.g., for selecting a book together with the names of its authors, use an “empty node”, denoted as `{}`, instead of the path concatenation `.` commonly used for combining path expressions.

The RQL path expression `{X}books:author.books:authorName{Y}` becomes in SeRQL `{X}<books:author>{<books:authorName>{Y}`. Furthermore, a number of short cuts addressing common queries are provided:

- Querying *multi-valued properties*. Properties in RDF can have multiple values and SeRQL provides a short cut to query several of these values in a single path expression. This allows for an easy formulation of Query 9 (note, the use of `<!` and `>` for enclosing a URI in contrast to the use of simple angle brackets (without the exclamation mark) for enclosing QNames)

```
CONSTRUCT      {X} <mybooks:co-author> {Y}
FROM           {Book} <books:author> {X, Y}
WHERE         X != Y
USING NAMESPACE books = <!http://example.org/books#>
              mybooks = <!http://example.org/books-rdfs-extension#>
```

This query is equivalent to

```
CONSTRUCT      {X} <mybooks:co-author> {Y}
FROM           {Book} <books:author> {X},
              {Book} <books:author> {Y}
WHERE         X != Y
USING NAMESPACE books = <!http://example.org/books#>
              mybooks = <!http://example.org/books-rdfs-extension#>
```

- Querying *multiple properties* of the same resource. Often one is interested not only in a single, but rather in multiple properties of a resource. The following query illustrates the syntactic short hand provided by SeRQL for this case: It selects the authors of all books with title "Bellum Civile" and a translator with name "J. M. Carter".

```
SELECT      Author
FROM        {Book} <books:title> {"Bellum Gallicum"};
           <books:translator> {<books:translatorName> {"J. M. Carter"};
           <books:author> {Author}
USING NAMESPACE books = <!http://example.org/books#>
```

Note, the use of the semicolon for separating expressions with the same subject.

- Querying *reified statements*. A reified expression can be queried by explicitly asking for the four triples associating the statement with its type (`rdf:Statement`), subject, predicate, and object. As this is often bothersome, SeRQL allows reified statements to be queried by enclosing the non-reified version of the statement in curly brackets.
- As all properties in RDF are essentially **optional** (unless considering schema languages such as OWL), queries where some information is retrieved if it is available without making the query fail if the information misses are a natural requirement for an RDF query language. SeRQL provides optional path expressions by enclosing an arbitrary path expressions (including those using the short hand constructs described above) in square brackets. E.g., the following query retrieves books together with their title and optionally their translators. If there is a translator and an age for that translator is specified, that information should also be returned.

```
SELECT      *
FROM        {Book} <books:title> {Title};
           [<books:translator> {Translator}]
```

```
[ <books:age> {Age} ] ]  
USING NAMESPACE books = <!http://example.org/books#>
```

Note, the nesting of the optional expressions and the use of * to include bindings for all variables from the query in the result.

Aside of these issues, it should be noted that the access to schema information has been reduced to the provision of both the original and the intransitive view of the two RDFS relations defining the subsumption hierarchy.

Easy of use. The expressed goal of the SeRQL authors has been a simplification of previous approaches such as RQL while retaining the, from the point of view of the SeRQL authors, most useful features that distinguish RQL from basic RDF query languages such as RDQL. Arguably, simply by reducing the number of language constructs, SeRQL is easier to grasp and use than RQL. Also, the syntactic short hands are, where possible, aligned with constructs from established RDF syntaxes such as N3 or N-Triples.

However, this is obviously paid for by reducing the expressiveness of the language in comparison to RQL, therefore requiring more effort for writing complex queries that could benefit from the more expressive features of RQL.

Functionality—Query Types. As expected, SeRQL can *not* express all of the queries from Section 3.2 that could be expressed in RQL but still provides more functionality than RDQL. Selection and extraction queries can be easily expressed in SeRQL with the same limitation as in the case of RQL, viz. that it is not possible to navigate arbitrary length paths in the graph, e.g., for returning all statements related to a resource or a “concise bounded description” as defined in [246].

In contrast to RQL, SeRQL currently neither provides set operations nor existential or all quantifiers. Therefore, Query 3 can not be expressed in SeRQL.

Thanks to the ability to construct new statements using the CONSTRUCT clause, SeRQL can express restructuring and simple inference queries as shown above. The restructuring Query 4 can be expressed as

```
CONSTRUCT      {Author} <mybooks:authored> {Book}  
FROM           {Book} <books:author> {Author}  
USING NAMESPACE books = <!http://example.org/books#>  
              mybooks = <!http://example.org/books-rdfs-extension#>
```

Aggregation queries can not be expressed, although [63] states that the addition of aggregation queries to SeRQL is planned.

As shown above, some simple inference queries such as Query 9 can actually be implemented in SeRQL and thanks to the RDFS-aware storage in Sesame the transitive closure of `rdfs:subClassOf` is provided in SeRQL. However, neither the transitive closure of arbitrary relations nor general recursion can be expressed.

Semantics. No formal semantics has been provided so far, however a formal algebraic model is being planned according to [63].

Complexity and implementation. No formal complexity study of the language has been published so far. There are currently two independent implementations of the language.

Reasoning. Again, only limited reasoning abilities are provided. Using the CONSTRUCT clause one can implement basic derivations as demonstrated above, however no recursion is provided.

Ontology awareness. SeRQL is RDFS aware in that it provides the ability to query both the explicitly stored subsumption relations and their transitive closure.

Project page:

Sesame <http://www.openrdf.org/>

Implementation(s):

available from the Sesame project page, an implementation in Prolog using the SWI-Prolog⁴ Semantic Web library is provided at <http://gollem.swi.psy.uva.nl/twiki/pl/bin/view/Library/SeRQL>

Online demonstration:

several ones (featuring not only SeRQL as query language, but also RDQL and RQL) accessible at <http://www.openrdf.org/demo.jsp>

4.3.3.3 eRQL

In contrast to SeRQL, which aims at providing a language more balanced between expressiveness and ease-of-use than RQL, eRQL [251] proposes a radical simplification using essentially a keyword-based interface similar to popular information retrieval systems. It is the expressed goal of the eRQL authors, to provide a “Google-like query language but also with the capacity to profit of the additional information given by the RDF data”⁵.

eRQL only provides three query constructs:

- *One-word queries.* Single keywords are valid eRQL queries. E.g., the query CAESAR returns all statements such that the string “CAESAR” occurs in their URI or literal value of the subject, predicate, or object of the statement using case-insensitive matching. Surprisingly, phrase queries (e.g., the phrase “Bellum civile”) do not seem to be expressible in eRQL.
- *Neighborhood queries.* Instead of returning only the statements directly containing a keyword as in the first case, neighborhood queries allow the user to also select all statements related to (i.e., “in the neighborhood of”) a such a statement. E.g., the query {{CAESAR}} returns all statements connected by at most two edges in the RDF graph to a node containing “CAESAR”. On the data from Figure 2.1, the following triples are returned:

```
_:1 books:author _:2.  
_:1 books:authorName "Julius Caesar".  
_:1 books:author _:3.  
_:1 books:title "'Bellum Civile'".  
_:1 books:translator _:4.
```

⁴ <http://www.swi-prolog.org/> ⁵ <http://www.dbis.informatik.uni-frankfurt.de/~tolle/RDF/eRQL/>

Using `{{{CAESAR}}}` would also include the names of the other others and the translator, as well as the triple classifying the book as `books:Historical_Essay`. eRQL allows any (finite) number of such brackets to select a neighborhood of the specified size around the base triples.

- *Conjunctive and disjunctive queries.* Both, neighborhood and one-word queries can be combined using the boolean operators AND and OR. No negation is provided, however.

eRQL does not allow the expression of most of the queries given in Section 3.2, since the abilities of eRQL are more akin to an information retrieval language than a conventional query language. However, to some extent Query 2 can be expressed simply by the query

```
{{"Bellum" AND "Civile"}}
```

This query returns all statements containing both the string “Bellum” and the string “Civile” in the URI or literal value of the subject, predicate, or object together with all statements reachable from these within two steps. However, this is only a vague approximation of the actual intent of the query. In particular, eRQL does not allow the selection of a neighborhood with previously unknown size around a resource (e.g., for obtaining a “concise-bounded descriptions” [246]). In contrast to the claims of the authors of eRQL, such limited neighborhood queries require a-priori knowledge of the schema of the data to be queried.

Nevertheless, eRQL is one of the few approaches aiming at a combination of information retrieval features and RDF querying, the need for which is evident when considering the use of RDF for improving searching in the (Semantic) Web.

Project page:

<http://www.dbis.informatik.uni-frankfurt.de/~tolle/RDF/eRQL/>

Implementation(s):

eRQLEngine, available from <http://www.dbis.informatik.uni-frankfurt.de/~tolle/RDF/eRQL/>

Online demonstration:

none

4.3.4 Query Languages using a Controlled Natural Language

4.3.4.1 Metalog

Metalog [182, 183, 181] is a system for querying and reasoning with Semantic Web data. It has been introduced around 1998 in [182] leading to the claim of the project page that “Metalog has been the first semantic web system to be designed, introducing reasoning within the Semantic Web infrastructure by adding the query/logical layer on top of RDF” [<http://www.w3.org/RDF/Metalog/>].

For two reasons, Metalog is notably different from most of the other RDF query languages discussed in this chapter:

- It combines querying with *reasoning* abilities such as implications and explicit representation of negative information.

- The language syntax is similar to a restricted form of natural language, where only certain keywords (and their order) is relevant and all other words are discarded. This allows for easy understanding of queries, although authoring still requires knowledge about the exact keywords and how to use them.

Metalog's syntax uses (English) sentences to describe a query. Within a sentence only three kinds of tokens are actually meaningful for the Metalog processor: variables (or "representations") are written in capital letters only, any quote-delimited string is recognized as a "name" and interpreted either as an RDF literal or a URI, and keywords connect variables and names to queries. The list of keywords available includes of statements, logical implication expressed by `then`, `imply`, or `implies`, definition of variables using `represents`, order for building RDF sequence containers, and various keywords for arithmetic expressions.

Metalog also provides a natural-language-like syntax for stating RDF triples.

The following Metalog program implements Query 1:

```
comment: some definitions of variables (or representations)
ESSAY represents the term "Essay"
      from the ontology "http://example.org/books#".
AUTHORED-BY represents the verb "author"
      from the ontology "http://example.org/books#".
IS represents the verb "rdf:type"
      from RDF "http://www.w3.org/1999/02/22-rdf-syntax-ns#".
BELLUM_CIVILE represents the book "Bellum_Civile"
      from the collection of books "http://example.org/books#".
comment: RDF triples written as Metalog statements.
BELLUM_CIVILE IS an ESSAY.
BELLUM_CIVILE is AUTHORED-BY "Julius Caesar".
BELLUM_CIVILE is AUTHORED-BY "Aulus Hirtius".
comment: a Metalog query
do you know SOMETHING that IS an ESSAY and that is AUTHORED-BY SOMEONE?
```

As answer to the query shown in the last line of the Metalog program, the interpreter answers with all definitions and the first two RDF triples as first result, the first and the last RDF triple as second result.

An interesting observation about Metalog can be drawn from [180]: A natural-language layer on top of the textual or XML syntax of any RDF query language might help, in particular non-experts, to quickly grasp the meaning of queries and what result to expect. In this sense, a translation of Metalog or a similar restricted form of natural language into some of the more traditional and, in many cases, better performing approaches discussed in this survey might be worth investigating.

Project page:

<http://www.w3.org/RDF/Metalog/>

Implementation(s):

prototype, available from the project page

Online demonstration:

none

4.3.5 Others

4.3.5.1 Algae

Algae⁶ [215, 217] is an RDF query language developed as part of the W3C Annotea project (<http://www.w3.org/2001/Annotea/>). The Annotea project provides a research platform for novel collaborative applications based on shared metadata such as Web annotations, bookmarks, comments, explanations, etc. When users access a Web site with an Annotea-enabled browser (such as Amaya (<http://www.w3.org/Amaya/>), Mozilla or Internet Explorer, the latter two require extensions for using Annotea), one or several *annotation servers* are contacted to deliver annotations for the currently visited Web site or to provide a classification of that Web site based on previously stored bookmarks.

The efficient retrieval of the related information from the annotation server is clearly of high relevance in such a setting. As the annotations are stored in RDF, an RDF query language, called Algae, has been developed to address the special needs of this application including the need for updates and simple reactive rules. In [217], Algae is described in more detail and proposed as a general-purpose RDF query language.

Algae is centered around two concepts:

- “Actions” are the directives `ask`, `assert`, and `fwrule` that determine whether an expression is used to query the RDF data, insert data into the graph, or to specify ECA-like rules. Only the first of these is mandatory for an Algae processor, the two others are defined in an extension module described in [216].
- Algae queries produce result sets containing not only bindings for query variables (as, e.g., RDQL [232]) does, but also triples from the RDF graph that constitute “proofs” for the solution, i.e., that are required to justify that a certain combination of bindings for the query variables is actually a match for the query. It is possible to combine several sub-queries in a single Algae expression in which case the results of each sub-query are combined into a single result set.

Syntactically Algae is based on N-triples (described in [123]) for representing and querying RDF triples. This simple triple syntax is extended by the above mentioned action directives and so-called “constraints”, written between curly brackets, that specify further arithmetic or string comparisons that must be fulfilled by a selected tuple.

To illustrate the abilities of Algae, consider Query 1. This query could be realized by the following Algae expression:

```
ns rdf = <http://www.w3.org/1999/02/22-rdf-syntax-ns#>
ns books = <http://example.org/books#>
read <http://example.org/books> ()
ask (   ?essay rdf:type      <http://example.org/books#Essay> .
       ?essay books:author  ?author .
       ?author books:authorName ?authorName )
collect( ?essay, ?author, ?authorName )
```

This query becomes more interesting, if we are only interested in the titles of essays written by “Julius Caesar” but also want the translators of such books returned, if there are any:

⁶ The current version is sometimes also referred to as Algae2, since there has been an earlier incarnation with more limited querying abilities. In this survey, we follow [217] in referring to the language as simply “Algae”.

Table 4.1 Algae result set

<code>?title</code>	<code>?translator</code>	<i>Proof</i>
"Bellum Civile"	"J. M. Carter"	<pre> _:1 rdf:type <http://exam...ks-rdfs#Essay>. _:1 books:author _:2. _:2 books:authorName "Julius Caesar". _:1 books:title "Bellum Civile". _:1 books:translator "J. M. Carter". </pre>

```

ns rdf = <http://www.w3.org/1999/02/22-rdf-syntax-ns#>
ns books = <http://example.org/books#>
read <http://example.org/books> ()
ask (
  ?essay rdf:type <http://example.org/books#Essay> .
  ?essay books:author ?author .
  ?author books:authorName "Julius Caesar" .
  ?essay books:title ?title .
  ~?essay books:translator ?translator .
)
collect( ?title, ?translatorName )

```

Note the use of ~ to declare the translator an optional triple in the query.

Such a query executed against the sample data shown in Figure 2.1, would return the following result set shown in Table 4.1.

Easy of use Although neither a visual syntax nor a natural language interface is provided, Algae queries are easy to write and comprehend, although the syntax does not reminiscent of any particular other query language. In particular, the choice of N-triples (described in [123]) for representing triples is helpful for easy query formulation. Some of the syntactic short hands, such as ~ require some getting used to, but offer the advantage of a concise query formulation.

Algae offers an extension mechanism that allows the basic functionality of the language, i.e. the ask action directive, to be separated from more advanced features such as updates and rules (with the action directives assert and fwrule) and sorting. A query can state which of the language extensions are required to evaluate a query.

Functionality—Query Types From the queries discussed in Section 3.2, Query 2 can not be expressed due to the lack of some form of closure or recursion. For the same reasons and additionally the lack of negation, Query 4 can not be expressed. The support for Queries 5 and 6 falls short due to the lack of aggregation operators. All other queries can be expressed in Algae, albeit most of them require the extended action directives discussed in [216].

For Query 4 one could use the following Algae query:

```

ns books = <http://example.org/books#>
read <http://example.org/books> ()
fwrule ask(
  ?book books:author ?author )
  assert ( ?author books:authored ?book )

```

Note, however that this query actually becomes an inference rule, as there are no means described to retract the old books:author relation.

Hence, this realization of Query 4 is more akin to a inference query such as Query 9 that could actually be implemented as

```
ns books = <http://example.org/books#>
read <http://example.org/books> ()
fwrule ask ( ?book books:author ?author1 .
             ?book books:author ?author2 { ?author1 != ?author2 }
            )
            assert( ?author1 books:co-author ?author2 )
```

Semantics No formal semantics has been published for Algae.

Complexity and implementation Algae has been implemented in the W3C Annotation Server as part of the Annotea Project. The data can be stored in-memory, in a relational database, as described in [218], or directly generated from application data. As discussed above, Algae provides extensions for updating and

No formal complexity study of the language has been published so far.

Reasoning In [216], an extension for Algae is described that allows the statement of rules for intensional data specification: If a given query succeeds, new data is added to the data store possible drawing data from the result of the query. These rules are more similar to ECA-rules, albeit with an action part limited to inserting new data, than to view definitions. No further description of reasoning mechanisms in Algae is provided.

Ontology awareness As already specified, the language ignores every kind of ontological aspect, including typing mechanisms. If such aspects must be considered, they must be treated like all the user-defined data.

It is, however, possible to implement the special semantics of RDFS relations such as `rdfs:subClassOf` using Algae rule notation. The following is an implementation of Query 8 in Algae, in that it adds the transitive closure over `rdfs:subClassOf` to the data store.

```
ns rdfs = <http://www.w3.org/2000/01/rdf-schema#>
read <http://example.org/books> ()
fwrule ask ( ?X rdfs:subClassOf ?Z.
            ?Z rdfs:subClassOf ?Y
            )
            assert ( ?X rdfs:subClassOf ?Y )
```

Project page:

<http://www.w3.org/2004/05/06-Algae/> and for the Annotea project <http://www.w3.org/2001/Annotea/>

Implementation(s):

W3C Annotation Server <http://annotest.w3.org/annotations>

Online demonstration:

Query interface to the W3C Annotation Server using Algae as query language: <http://annotest.w3.org/annotations?explain=false>

4.3.5.2 iTQL

For the Kowari Metastore, an open source, scalable, transaction safe database for the storage of metadata, an RDF query language called iTQL [1] has been defined. iTQL provides commands not only for querying (`select`), but also for updates (`delete`, `insert`) and transaction management (`commit`, `rollback`). The syntax of iTQL is similar to SQL and therefore also reminds of RDQL. As for RDQL, the querying abilities of the language are rather limited, mostly simple selection is supported.

To illustrate the abilities of the query language, consider again Query 1 and a possible realization in iTQL:

```
alias <http://example.org/books#> as books;
alias <http://www.w3.org/2000/01/rdf-schema#> as rdfs;
alias <http://www.w3.org/1999/02/22-rdf-syntax-ns#> as rdf;
select $essay, $author, $authorName
where $essay <books:author> $author
and $author <books:authorName> $authorName
and $essay <rdf:type> $type
and (trans($type <rdfs:subClassOf> <books:Essay>)
or $type <tk:is> <books:Essay>)
```

As illustrated in the query, iTQL provides the function `trans` as means for computing transitive closure of a relation (such as `rdfs:subClassOf`) and therefore also resources not directly classified as `books:Essay` but rather as one of its subclasses are returned. Paths of arbitrary length in the graph can be traversed using another special function called `walk`. The above query could also be expressed using `walk`.

Worth mentioning is the ability to sort the resulting answers and to provide access to answers in a paged mode using `limit` and `offset` as in SQL.

Also, in contrast to the SquishQL-family of query languages discussed in Section 4.3.1, iTQL allows the specification of nested queries.

Project page:

<http://www.kowari.org>

Implementation(s):

production use implementation as part of the Kowari Metastore and used in the commercial product Tucana Knowledge Server

Online demonstration:

none

4.3.5.3 N3QL

A restricted subset of Notation 3 [32] (short N3), an alternative syntax and extension for RDF that introduces rules, variables, and quoting for easy expression of statements about statements.

Although, as noted in [36], the rules mechanism provided in N3 allows for similar capabilities as those expected from a query language, [36] proposes a syntax for a query language using more conventional means such as `select-where` clauses.

The essential difference between N3QL and most of the other query languages for RDF discussed in this survey, is that a query is an N3 expression and all “keywords” in the query

are actually RDF properties of an RDF node representing the query (usually a blank node, but it is also allowed to assign identifiers to queries). To illustrate this difference, consider Query 1 and its realization in N3QL:

```
@prefix books: <http://example.org/books#>.
@prefix n3ql: <http://www.w3.org/2004/q1#>.
@prefix rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#>.
[] n3ql:select { n3ql:result n3ql:is (?book ?author ?authorName) };
   n3ql:where { ?book rdf:type books:Essay;
                ?book books:author ?author;
                ?author books:authorName ?authorName }.
```

The result of such a query is the RDF graph specified in the `n3ql:select` clause, in this case a set of RDF collections (indicated by the collection constructor `()`) containing a binding for each of the three selected variables.

[36] seems to indicate that the semantics of such a query is equivalent to the semantics of a rule where the `where` part of the query is the premise of the rule and the `select` part the implication. However, N3 rules can be used to implement, e.g., the transitive closure of an RDF property:

```
@prefix rdfs: <http://www.w3.org/2000/01/rdf-schema#>.
{?x rdfs:subClassOf ?z; ?z rdfs:subClassOf ?y}
=> {?x rdfs:subClassOf ?y}
```

Such an ability is not attributed to queries written in the above syntax.

Since the description of N3QL is very sketchy, it is hard to evaluate its expressive power: In particular, it is not clear which of the syntactic constructs of N3 can be used in N3QL. [36] states that N3QL is a restricted form of N3 where formulae cannot be nested and literals cannot be subjects of statements.

In particular, the N3 syntax for anonymous nodes, for navigating in the RDF graph using path expressions, and for quantifying variables is assumed to be available in N3QL too. This allows for very concise formulation of queries such as “Return all books written by an author with name ‘Julius Caesar.’”:

```
@prefix books: <http://example.org/books#>.
@prefix n3ql: <http://www.w3.org/2004/q1#>.
[] n3ql:select { n3ql:result n3ql:is (?book) };
   n3ql:where { ?book!books:author!books:authorName ‘Julius Caesar’ }.
```

Project page:

<http://www.w3.org/DesignIssues/N3QL.html>

Implementation(s):

prototype, CWM <http://www.w3.org/2000/10/swap/doc/cwm.html>, EulerSharp <http://eulersharp.sourceforge.net/2003/03swap/>

Online demonstration:

none

4.3.5.4 PerlRDF Query Language

PerlRDF⁷ is a collection of APIs for parsing, storing, and querying RDF developed by Ginger Alliance⁸. As part of this project, also an RDF query language has been specified. The details of this language are described in [11].

The query language provides path expressions similar to RQL's general path expression and uses a familiar `Select...From...Where` syntax. An implementation of Query 1 modified to return only those books with a translator named "J. M. Carter" in this query language is shown in the following:

```
Select  ?book, ?author, ?author->books:authorName
From    books:Essay::?book->books:author{?author}
Where   ?book->books:translator->books:translatorName=>'J. M. Carter'
Use books For [http://example.org/books#]
```

The query has deliberately written to demonstrate different features of the query language, e.g., the ability to use path expression not only in the `From`-clause (as, e.g., in RQL), but also in the `Where`- and even in the `Select`-clause of the query. Furthermore, syntactic short hands for specifying the type of the `?book` variable and the constraint on the literal value selected by the path in the `Where` clause are used. This query is equivalent to the following one, that does not use any of these advanced constructs

```
Select  ?book, ?author, ?authorName
From    ?book->books:author{?author}->books:authorName{?authorName},
        ?book->rdf:type{books:Essay},
        ?book->books:translator->books:translatorName{?translatorName}
Where   ?translatorName = 'J. M. Carter'
Use books For [http://example.org/books#]
      rdf For [http://www.w3.org/1999/02/22-rdf-syntax-ns#]
```

The most interesting syntactic short hands are the use of `books:Essay::?book` to type the `?book` variable using a syntax reminiscent of the underlying Perl language and the use of the `=>` operator to mark the so-called "target element" of a path and place a restriction on the (literal) value of that element.

Project page:

http://www.gingerall.com/charlie/ga/xml/p_rdf.xml

Implementation(s):

available from the project page

Online demonstration:

http://rdf-demo.gingerall.cz/charlie/rdf/act/rdf_demo.act

4.3.5.5 R-DEVICE Deductive Language

The R-DEVICE system, presented in [23], is a "deductive object-oriented knowledge-base system for querying and reasoning about RDF metadata" [<http://lpis.csd.auth.gr/systems/r-device.html>]. It is a reimplementaion of the X-DEVICE language [22] in the C Language Integrated Production System (CLIPS, see <http://www.ghg.net/clips/CLIPS.html>) using the CLIPS Object-Oriented Language (COOL).

The mapping from RDF triples to objects is achieved in the following way:

⁷ http://www.gingerall.com/charlie/ga/xml/p_rdf.xml ⁸ <http://www.gingerall.com/>

- All resources are represented as objects where the type is determined by the `rdf:type` property of the resource. For resources that are classified in multiple classes a dummy class represents a common subclass of all the classes the resource is classified in.
- Properties are realized as multi-slots (slots with multiple values) in the class that is the domain of the property. If no domain is given the property can be applied to any resources, therefore is added as a slot to the class representing `rdfs:Resource` (the top of the resource object hierarchy).

New assertions (generated, e.g., through rules) can require dynamic class and/or object re-definitions.

To illustrate the syntax of R-DEVICE consider the following implementation of Query 1:

```
(deductiverule q1
  ?book <- (? (rdf:type books:Essay) (books:author ?author))
  ?author <- (? (books:authorName ?authorName))
=>
  (result (book ?book) (author ?author) (authorName ?authorName))
)
```

Note, the production-rule like syntax of R-DEVICE. R-DEVICE also provides constructs for traversing arbitrary length paths of slots and objects (properties and resources) both with and without restriction on the type of slot that may be traversed. This allows to implement both Query 2, where we want to collect all things related to the book with title “Bellum Civile” (this is indicated in R-DEVICE by an unconstrained path of arbitrary length from the book `?book` to the related resources `?related`),

```
(deductive rule q2
  ?book <- (? (rdf:type books:Essay) (books:title ‘Bellum Civile’) ((?p) ?related)
=>
  (result (book ?book) (related ?related))
)
```

and Query 8, where the transitive closure of the `rdfs:subClassOf` relation is to be computed. The latter query can be expressed using a recursive sub-path `rdfs:subClassOf`.

Project page:

<http://lpis.csd.auth.gr/systems/r-device.html>

Implementation(s):

available from project page

Online demonstration:

none

4.3.5.6 RDF-QBE

In [223] a language for querying RDF graphs following the well-established “query by example” paradigm [260, 261] is proposed. Essentially, an RDF graph (described in Notation 3 syntax [32]) is used to describe the query pattern that should be found in the data. Variables in the pattern are expressed as blank nodes without explicit node identifiers as described in [156].

This leads to a major restriction of the approach: query patterns may only form a tree not a graph.⁹

Query 1 can be expressed in RDF-QBE as

```
[ ] a books:Essay; books:author [ books:authorName [ ] ] .
```

When considering what this query should return, the problem of handling blank (or anonymous) nodes in RDF must be addressed. Either some identifier (scoped only within the collection of data considered at the moment, for more details see [156]) is assigned to the blank node or the blank node is represented as the collection of its properties. This issue actually has to be considered by all RDF query languages, in particular as assigning an identifier to blank nodes does not mean that blank nodes can be treated as a node with explicit identifier.

In general, RDF-QBE provides a very convenient, easy to read syntax, but the trade-off (acknowledged in [223]) is the rather low expressiveness of the language, as detailed in Appendix B.

Project page:

none

Implementation(s):

described in [223], but not publicly available

Online demonstration:

none

4.3.5.7 RDFQL

RDF Gateway [2] is a platform for developing and deploying Semantic Web applications combining a “native” RDF database engine with a Web server and a server-side scripting language. The RDF database engine allows for the integration of standard and Semantic Web using so-called “virtual tables” and inference rules for deductive reasoning (so far, libraries for OWL and RDFS are provided). A graphical editor for RDF graphs statements is provided for easy creation of new data. To enable basic interoperability with different Semantic Web tools, several RDF serialization formats are supported, viz. RDF/XML, N3 and NTriples (cf. Appendix A).

The RDF Gateway uses a proprietary query language, referred to as RDFQL, described in [3]. Although in many ways similar to RDQL there are several noteworthy differences:

- Transaction management is realized in RDFQL by database commands for starting and committing or undoing (rollback) of a transaction.
- SQL-like update commands are provided, including a full data definition language.
- Data can be stored in data sources (often referred to as “tables”, although they differ from tables in SQL database by having a fixed schema as they are only meant to store RDF triples or quads) that can be either disk-based, in-memory or external data sources identified, e.g., by an URI.

⁹ Contrary to the claim in [223], this does however not reduce the problem to tree matching, as the data is still graph shaped.

- Using the command INFER, deductive rules can be defined as part of a RULEBASE to be used when querying. This allows, e.g., to specify the semantics of RDFS in the following way (note, that an RDF statement with subject *S*, predicate *P*, and object *O* is written in RDFQL as {?P ?S ?O}, i.e., in prefix notation; note also the use of uri(?u)=?u to detect whether the object of an RDF statement is a resource (in which case it has an URI and that URI is equal to the “value” of the resource itself) or a literal):

```

RULEBASE rdfs
{
  INFER {[rdf:type] ?a [rdf:Property]} from {?a ?x ?y};
  INFER {[rdf:type] ?x ?z} from {[rdfs:domain] ?a ?z} and {?a ?x ?y};
  INFER {[rdf:type] ?u ?z} from {[rdfs:range] ?a ?z}
    and {?a ?x ?u} and uri(?u)=?u;
  INFER {[rdf:type] ?x [rdfs:Resource]} from {?a ?x ?y};
  INFER {[rdf:type] ?u [rdfs:Resource]} from {?a ?x ?u} and uri(?u)=?u
  INFER {[rdfs:subPropertyOf] ?a ?c}
    from {[rdfs:subPropertyOf] ?a ?b} and {[rdfs:subPropertyOf] ?b ?c}
  INFER {?b ?x ?y} from {[rdfs:subPropertyOf] ?a ?b}
    and {?a ?x ?y}
  INFER {[rdfs:subClassOf] ?x [rdfs:Resource]}
    from {[rdf:type] ?x [rdfs:Class]}
  INFER {[rdfs:subClassOf] ?x ?z} from {[rdfs:subClassOf] ?x ?y}
    and {[rdfs:subClassOf] ?y ?z}
  INFER {[rdf:type] ?a ?y} from {[rdfs:subClassOf] ?x ?y}
    and {[rdf:type] ?a ?x}
}

```

Query 1 can be implemented in RDFQL by the following program:

```

session.namespaces["books"] = "http://example.org/books#";
var booksdata = new DataSource("http://example.org/books");
SELECT ?essay, ?author, ?authorName USING booksdata WHERE
  {[rdf:type] ?essay [books:Essay]}
  and {[books:author] ?essay ?author}
  and {[books:authorName] ?author ?authorName}
ORDER BY ?authorName DESC;

```

Note again the (in the RDF context) uncommon notation of statements with predicate first. For illustration of the ability of RDFQL to return ordered result, an ORDER BY clause is added to the query that orders the result by the name of the author. Using a rule base as the one shown above for RDFS, even resources classified by a sub-class of books:Essay will be returned.

Project page:

<http://www.intellidimension.com/>

Implementation(s):

within the RDF Gateway, personal edition (limited to non-commercial use and the number of connections allowed) available from project page

Online demonstration:

none (however, the project page can serve as a show case as it is implemented using RDF Gateway)

4.3.5.8 TRIPLE

TRIPLE [236, 237, 134] is a rule-based query, inference, and transformation language for RDF data based upon ideas published in [90] and with a syntax close to F-Logic [152]. The use of F-Logic for querying semi-structured data such as XML or RDF is natural, as one of the strengths of F-Logic approaches is the ability to handle data without a fixed schema demonstrated, e.g., in [171]. Other F-Logic based approaches are, e.g., XPathLog (see Section 4.2.1.7) and the commercial ontology management platform Ontobroker¹⁰.

TRIPLE has been designed to address in particular two weaknesses of previous approaches for RDF query languages:

- Most previous proposals provide a number of predefined constructs implementing the specific semantics of, e.g., RDFS or OWL. The disadvantage of such an approach is the lack of extensibility, although extensibility is a crucial feature of the underlying representation formalism RDF. In contrast, TRIPLE only offers a basic, rule-based language for Horn logic, that is in large parts identical to F-Logic [152]. This language can be used, where possible, to implement the semantics of, e.g., RDFS. Where Horn logic is not sufficient, as, e.g., in the case of OWL, TRIPLE is designed to be extended by external modules implementing, e.g., an OWL reasoner.

Building upon [159, 160] for expressing Topic Maps in RDF and [192] for representing UML in RDF, the authors argue in [237] that TRIPLE could also be used to query other representation formalisms for metadata other than RDF.

However, this claim is not demonstrated and, although recent work [116] gives some impression of an integrated query language for RDF and Topic Maps, the adequacy of TRIPLE for querying Topic Maps is questionable in light of the rather awkward mappings from Topic Maps to RDF proposed so far (e.g., many of these approaches result almost exclusively in reified statements).

- Due to the foundation in Horn logic, TRIPLE provides not only a well-defined semantics but also fairly powerful reasoning capabilities (that can be further enhanced by extension modules), both in contrast to previous approaches. In particular, the use of a rule language for both querying and reasoning about the queried data is a natural choice, even more so in the context of the Semantic Web.

Following [237] one can identify a number of areas where TRIPLE differs from basic Horn logic (and logic programming languages such as Prolog). Most of these differences are either related to specificities of the RDF data model or the choice of representing properties similar to slots in F-Logic:

- **Identifying resources:** Resources are identified in RDF by URIs. TRIPLE supports both namespaces and general resources abbreviations (e.g., `isa := rdfs:subClassOf` to simplify the notation of URIs. TRIPLE assumes that all resources are identified properly by an URI, anonymous resources are not considered so far (there is some indication that this will change in the future in [134]).
- **Representing and querying statements:** RDF statements are represented as slots and slot values of the subject in a statement, i.e., `subject[predicate -> object]`. This

¹⁰ <http://www.ontoprise.de/products/ontobroker>

allows for easy grouping and nesting of statements. As in F-Logic, Path expressions inspired by [107] can be used to traverse several predicates at once.

- **Reification:** In contrast to many other RDF query languages, TRIPLE provides concise support for reified statements by enclosing such statements in angle brackets, e.g., `Julius_Caesar[believes-><Junius_Brutus[friend-of -> Julius_Caesar]>]`.
- **Explicit model specification:** Similar to the module syntax in some F-Logic systems, TRIPLE allows the explicit specification of the model in which a statement or atom is true. The model is again identified by an URI and appended to the statement or atom by `@`.

Finally, one should note that TRIPLE differs from common logic programming languages such as Prolog in requiring all variable to be explicitly quantified.

With this syntax RDF statements and queries can be expressed in TRIPLE. Assuming the data from Figure 2.1 has been loaded as part of a model identified by `http://example.org/books` the following TRIPLE program implements Query 1:

```

rdf      := 'http://www.w3.org/1999/02/22-rdf-syntax-ns#'.
books    := 'http://example.org/books#'.
booksModel := 'http://example.org/books'.
FORALL B, A, AN result(B, A, AN) <-
    B[rdf:type -> books:Essay;
     books:author -> A[books:authorName -> AN]]@booksModel.

```

Note, the use of both nesting and grouping of statements for this query. In this formulation, this query selects only resources directly classified as `books:Essay`, below it is discussed how this query can be modified to properly select all resources classified as `books:Essay` or any of its sub-classes in the RDFS subsumption hierarchy.

As discussed above, the specific semantics of different RDF vocabularies such as RDFS or OWL is provided on top of the basic language layer, either as external modules or implemented using the Horn logic reasoning provided by TRIPLE. In [237] the following implementation of RDFS in TRIPLE is given:

```

rdf      := 'http://www.w3.org/1999/02/22-rdf-syntax-ns#'.
rdfs     := 'http://www.w3.org/2000/01/rdf-schema#'.
type     := rdf:type.
subPropertyOf := rdfs:subPropertyOf.
subClassOf   := rdfs:subClassOf.
FORALL Md1 @rdfschema(Md1) {
    transitive(subPropertyOf).
    transitive(subClassOf).
    FORALL O,P,V O[P->V] <-
        O[P->V]@Md1.
    FORALL O,P,V O[P->V] <-
        EXISTS S S[subPropertyOf->P] AND O[S->V].
    FORALL O,P,V O[P->V] <-
        transitive(P) AND EXISTS W (O[P->W] AND W[P->V]).
    FORALL O,T O[type->T] <-
        EXISTS S (S[subClassOf->T] AND O[type->S]).
}

```

Note, that this implements only part of the RDFS semantics, e.g., inference from range and domain restrictions of properties is not provided. However, this is not due to a limitation of TRIPLE, as adding the following two rules completes the implementation of RDFS in TRIPLE.

```

FORALL S,T S[type->T] <-
  EXISTS P, O (S[P->O] AND P[rdfs:domain->T]).
FORALL O,T O[type->T] <-
  EXISTS P, S (S[P->O] AND P[rdfs:range->T]).

```

With these rules, the implementation of Query 1 in TRIPLE shown above only needs to be modified with respect to the model it is evaluated against: instead of `@booksModel`, `@rdfschema(booksModel)` is used, that is the original model “expanded” with the above rules implementing RDFS semantics.

In [237], not only the textual syntax for TRIPLE discussed so far is proposed but also an RDF (and therefore) XML representation of the language itself (more precisely, for the basic language layer).

Easy of use. Although TRIPLE employs a syntax less familiar to the average query programmer than languages inspired by main-stream query languages such as SQL or XPath, the close alignment with F-Logic allows users knowledgeable in logic programming languages to become acquainted with the language quickly. However, the lack of a visual syntax or a natural language interface makes it hard for non-experts to formulate. An RDF (and therefore XML) syntax is provided that can indeed be queried by TRIPLE itself.

As discussed above, TRIPLE allows the explicit specification of the model(s) for which a statement or formula should hold. Even more, reasoning methods beyond the Horn logic formulae provided by the basic TRIPLE language can be implemented by external modules.

The most striking weakness of the TRIPLE language is a direct consequence of the generality claimed in [237]. Although data represented in very different formalisms such as RDF, Topic Maps, or UML can be queried in TRIPLE (e.g., by translating the data to RDF), this leads to rather awkward representations of many language features. Even for RDF there are certain aspects of RDF, viz. containers, collections, and anonymous nodes, that are not considered by TRIPLE and can not easily be added to the language.

Functionality—Query Types. As just discussed, TRIPLE’s generality is in some cases paid for by a lack of adequacy for the data representation formalisms claimed to be supported. Regarding the different query classes discussed in Section 3.2, one can observe, that most of them can be expressed in TRIPLE as already demonstrated for Query 1 and 8. However, the language does not support aggregation.

Semantics. In [237] the semantics of TRIPLE is given by a mapping of the TRIPLE-specific features to standard Horn logic expressions.

Complexity and implementation. No formal complexity study of the language has been published. So far, only a single prototype implementation of TRIPLE is available.

Reasoning. As discussed above, TRIPLE not only offers full Horn logic reasoning as part of the basic language but is designed to allow the extensions with specific reasoners, e.g., for handling OWL ontologies.

Ontology awareness. Although the basic language does not provide any ontology-specific abilities, ontology awareness can be obtained either by implementing the semantics of the ontology language in the basic TRIPLE Horn logic (as demonstrated for RDFS above) or by extending the language with specific reasoners tailored for the needs of the ontology language such as OWL.

Project page:

<http://triple.semanticweb.org/>

Implementation(s):

available from the project page

Online demonstration:

<http://ontoagents.stanford.edu:8080/triple/> (not functional at the time of writing), some information about projects realized with TRIPLE demonstrating its abilities are available from the project page

4.3.5.9 WQL

Ivanhoe [164] is a frame-based API following ideas from [142, 161] for the Nokia Wilbur Toolkit [166], a collection of APIs for processing XML, RDF, and DAML written in CLOS (Common Lisp Object System) and introduced in [162]. In Ivanhoe, resources described using RDF and/or DAML are represented as frames with a slot for each property of the resource. The (possibly multiple) values of a slot correspond to objects of RDF statements with the resources represented by the frame as subjects. In [165], a comparison of a subset of the Ivanhoe API, referred to as Wilbur Query Language or WQL, is described along the criteria from [131]: Three variants of WQL are discussed:

- the basic path-based query language that allows the selection of some resource reachable by a path via `value`, the selection of all resources reachable by a path via `all-values`, and the test whether two resources are connected by a certain path using `relatedp`;
- the embedding of that query language into Common Lisp (abbreviated as WQL+CL). It is not clear, if there is a restriction on the allowed Common Lisp expressions for this language layer.
- based upon the “transparent” (or “hidden”) inference extensions for Wilbur presented in [163], the final layer, abbreviated WQL+CL+inference, uses WQL+CL as query language but assumes a data store providing inferencing, e.g., for implementing RDFS.

For this survey we concentrate on the first query language WQL only, as WQL+CL is more akin to a programming language than a query language. However, where appropriate the “transparent inferencing” provided by Wilbur is considered when evaluating WQL.

Consider a query that returns the labels of all classes the book with identified by `http://example.org/books#Bellum_Civile` is classified under on the data from Figure 2.1:

```
(setf *db* (make-instance 'db))
(load-db (make-url "http://example.org/books")
         :locator "http://example.org/books")
(add-namespace "books" "http://example.org/books#")
(all-values !"http://example.org/books#Bellum_Civile"
 '(:seq !rdf:type (:seq (:rep* !rdfs:subClassOf) !rdfs:label)))
```

Note, the use of the `:seq` operator for constructing a sequence of slots (or relations in RDF terms) to be traversed by the query and of the `:rep*` operator for traversing the transitive closure of a slot/relation. `all-values` returns all resources (represented as frames) reachable by the specified path from the source frame, i.e., the frame identified by `http://example.org/books#Bellum_Civile`.

Project page:

Wilbur Toolkit: <http://wilbur-rdf.sourceforge.net/>

Implementation(s):

available from the project page

Online demonstration:

none

4.4 Topic Maps Query Languages

4.4.1 tolog: Logic Programming for Topic Maps

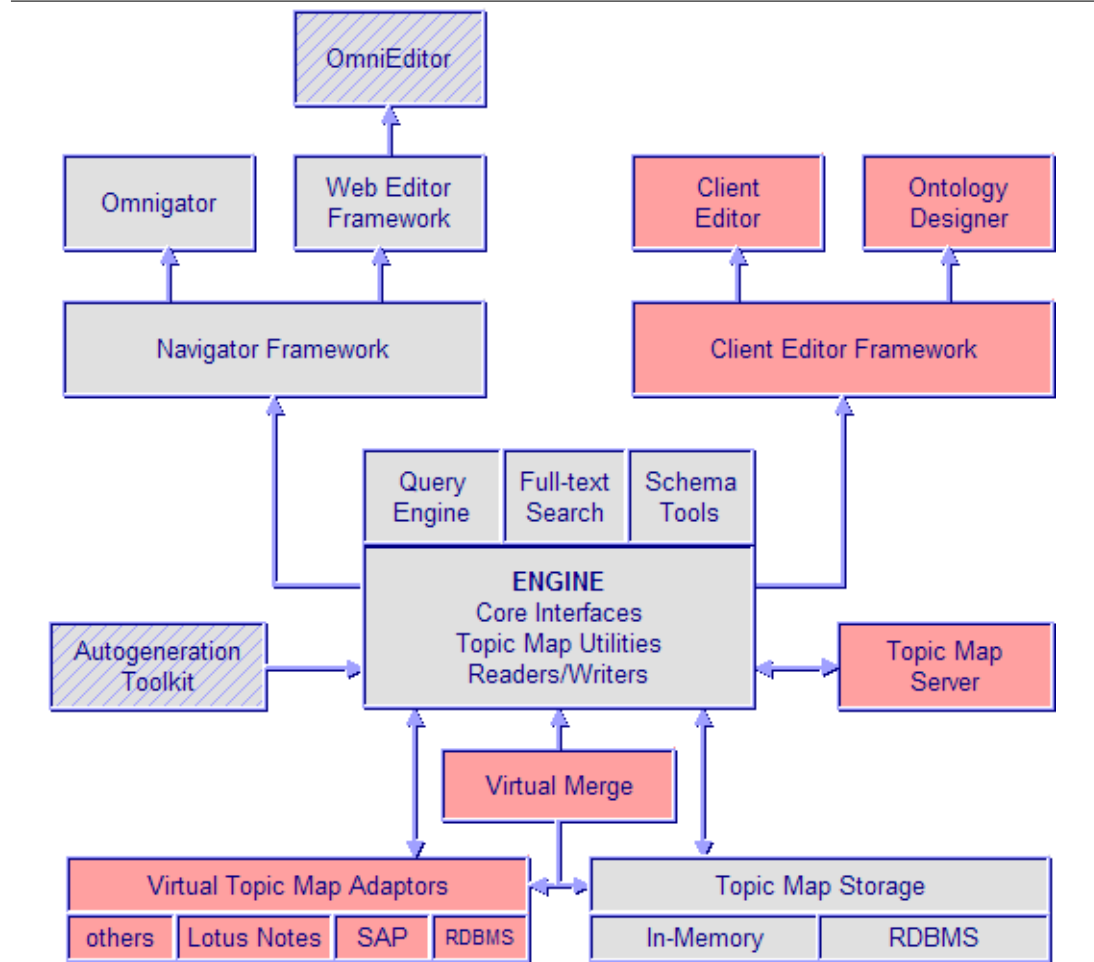
tolog [113] has been developed as part of the Ontopia Knowledge Suite¹¹, providing access to the core query engine (cf. Figure 4.2). It has also recently (April 2004) been selected as an initial straw-man for the ISO Topic Maps Query Language currently under development. The language is specified in [113] (incomplete), a gentle introduction is presented in the tolog tutorial [114], and in [111, 112] language design and evolution are addressed.

The design is most notably influenced by logic programming languages such as Prolog and logic-based query languages such as Datalog. Some of the more common syntactic constructs have been adapted to an SQL-like feeling to appeal to query authors without background in logic programming. The influence from logic programming is obvious when considering the basic query constructs used in tolog (here and in the following, this discussion of tolog is mostly oriented at [114], since the specification is in many points still unfinished; where possible, extensions described in [112] have been considered):

- **Identifiers.** tolog provides several means for identifying a Topic Maps construct, most notably based on the (internal) ID of a topic and its subject indicator. E.g. the topic (type) “Novel” in the sample data could be addressed either using its ID, i.e., `Novel`, or its subject indicator `i"http://example.org/books#Novel"`. Note the `i` prefix in the latter case to distinguish between different identifiers used in a topic map. As usual, URI prefixes can be used to abbreviate such expressions, e.g., under the prefix definition using `books` for `i"http://example.org/books#"` one can write simply `books:Novel` to address that topic by its subject indicator. Note, that these prefixes are *not* equivalent to (XML) namespaces, as also the indicator which of the topic identifiers to use is included.
- **Variables.** In contrast to Datalog or Prolog, variables are prefixed with a `$` wherever they are referenced (this is both for simplicity and to be able to allow uppercase topic IDs). Per default all variables occurring in a query are returned, however using `select $var1, $var2, ... from query` a projection on the variables `$var1, $var2, ...` can be obtained.

¹¹ <http://www.ontopia.net/solutions/products.html>

Figure 4.2 Overview of the Ontopia Knowledge Suite
 (from <http://www.ontopia.net/solutions/products.html>, © Ontopia)



- **Predicates.** The original tolog proposal (cf. [113]) provides for two kinds of predicates: built-in and dynamic association predicates. For each association type occurring in a topic map, there is a so-called dynamic association predicate that allows querying the extent of the association type. E.g., to query the authors of book **b1** in the sample topic map shown in Chapter 2, one can write `authors-for-book(b1, $AUTHOR: author)`. Note, the use of the association role to identify which of the two associated topic is the author. In this case, the query processor might be able to infer the type of `$AUTHOR` from the type of **b1** and the fact that there are only two topics involved in an `authors-for-book` association. However, in a query such as `authors-for-book($BOOK, $AUTHOR)` the need for specifying the association roles is obvious. Analogously to dynamic *association* predicates there are also dynamic *occurrence* predicates.

The only two built-in predicates are `instance-of($INSTANCE, $CLASS)` and `direct-instance-of($INSTANCE, $CLASS)`, i.e., predicates that implement the special semantics of the subsumption hierarchy.

Although the dynamic association and occurrence predicates allow for easy authoring of queries, they require that the query author is aware of the ontology of the data to be queried. This disadvantage has been addressed in a recent proposal on extending tolog [112] and already incorporated into the language tutorial [114]: a number of additional built-in predicates for enumerating the associations, association roles, occurrences, and topics are provided, that allow querying arbitrary topic maps without a-priori knowledge of the types used in the topic maps. E.g., Query 2 can only be implemented using these predicates:

```
select $RELATED from
  title($BOOK, "Bellum Civile"),
  related($BOOK, $RELATED)?
related($X, $Y) :- {
  role-player($R1, $X), association-role($A, $R1),
  association-role($A, $R2), role-player($R2, $Y) |
  related($X, $Z), related($Z, $Y)
}.
```

This query also demonstrates the use of conjunctions (denoted by a comma as in prolog), disjunctions (denoted by an expression parenthesized with curly brackets in which the disjuncts are separated by `|`), and inference rules discussed below in detail. Note, the use of the built-in predicates `role-player` and `association-role` that related association roles with topics, respectively associations with association rules.

Using the notation for disjunctions, one can also specify optional queries.

- **Inference.** As tolog is based on Datalog, it uses a similar syntax for inference rules. E.g., the built-in predicates `instance-of` and `direct-instance-of` can indeed be implemented using only dynamic association predicates and inference rules as follows (cf. [113]):

```
direct-instance-of($INSTANCE, $CLASS) :-
  i"http://psi.topicmaps.org/sam/1.0/#type-instance"(
    $INSTANCE : i"http://psi.topicmaps.org/sam/1.0/#instance",
    $CLASS : i"http://psi.topicmaps.org/sam/1.0/#class").
super-sub($SUB, $SUPER) :-
  i"http://www.topicmaps.org/xtm/1.0/core.xtm#superclass-subclass"($SUB, $SUPER).
```

```

    $SUB   : i"http://www.topicmaps.org/xtm/1.0/core.xtm#subclass",
    $SUPER : i"http://www.topicmaps.org/xtm/1.0/core.xtm#superclass").
descendant-of($DESC, $ANC) :- {
  super-sub($DESC, $ANC) |
  super-sub($DESC, $INT), descendant-of($INT, $ANC)
}.
instance-of($INSTANCE, $CLASS) :- {
  direct-instance-of($INSTANCE, $CLASS) |
  direct-instance-of($INSTANCE, $DCLASS), descendant-of($DCLASS, $CLASS)
}.

```

Note, the use of the subject indicator to access the standardized type-instance and superclass-subclass associations.

Inference rules can also use negation, however according to [112] the semantics of negation in tolog is not yet fully specified.

Aside of these central concepts, tolog also provides constructs for aggregation and sorting ([112] mentions the need for additional aggregation functions), paged queries using `limit` and `offset` clauses similar to SQL, and means for defining and using modules of, e.g., inference rules. Furthermore, recent versions of tolog offer initial support for function libraries on simple data types similar to [178].

Easy of use. The textual syntax of tolog is closely aligned with logic programming languages in the style of Prolog or Datalog. Most of the Topic Maps extensions are rather straight-forward. However, neither a visual syntax nor a natural language interface are provided.

Functionality—Query Types.

Semantics. No formal semantics has been published.

Complexity and implementation. No formal complexity study of the language has been published. Tolog has been implemented in the Ontopia Knowledge Suite¹² and in the open source Topic Maps toolkit TM4J¹³. With tolog selected as initial straw-man for the upcoming ISO Topic Maps Query Language, a more wide-spread adoption of the language is to be expected.

Reasoning. Roughly the same reasoning abilities as in Prolog are provided, although the handling of negation is unclear in the current documents describing tolog.

Ontology awareness. tolog can access and query type information included in a topic map. In particular, the special semantics of the associations defining the subsumption hierarchy is considered. No further support for ontology languages is provided. However, it is likely that the ISO Topic Maps Constraint Language¹⁴ currently under development will be supported in the future.

Project page:

<http://www.ontopia.net/omnigator/docs/query/tutorial.html> (language tutorial)

¹² <http://www.ontopia.net/solutions/products.html>

¹³ <http://tm4j.org/>

¹⁴ <http://www.isotopicmaps.org/tmcl/>

Implementation(s):

as part of the Ontopia Knowledge Suite (<http://www.ontopia.net/solutions/products.html>) and the open source Topic Maps toolkit TM4J (<http://tm4j.org/>).

Online demonstration:

Omnigator is a show-case application implemented using the Ontopia Knowledge suite. It is available from <http://www.ontopia.net/> and an online demonstration can be accessed at <http://www.ontopia.net/omnigator/models/index.jsp> (however, there seems to be no way to directly test tolog queries in the online demonstrator).

4.4.2 AsTMA?: Functional-style Querying of Topic Maps?

AsTma? is an “experimental” Topic Maps query language integrated with the rest of the AsTMA language family developed at the Australian Bond University. It is defined in [16]. A language tutorial is also provided [17].

AsTma? has a rather different flavor compared to other Topic Maps languages as it is most similar to functional XML query languages in the style of XQuery [41]. It specifies several different path languages than can be used for accessing data in topic maps. These data can then be further processed by various query constructs, in particular for constructing XML output.

Query 1 can be implemented by the following AsTma? query:

```
<books>
{
  forall [$book (Writing)] in http://example.org/books
  return
  <book>
  { $book,
    forall $author in ($book -> author / author-for-book) return
    <author>
    { $author }
    <name>{ $author/bn}</name>
  }
  </author>
}
</book>
}
</books>
```

Note the almost identical syntax in comparison to XQuery. The query first selects all topics with topic type `Writing` in `http://example.org/books`. For each of these books, the author is queried by traversing the `author-for-book` association and selecting the topics with `author` role. Finally, the basename of each author is returned by the expression `$author/bn`.

One of the interesting features of AsTma? is the great variety offered for accessing topics and associations: based on a path expression in one of the path languages or on a constraint written in the AsTMA!. This allows the above query to be formulated as

```
<books>
{
  forall [$book (Writing)] in http://example.org/books
  return
  <book>
  { $book,
    forall [ (author-for-book)
              Writing : $book
              author: $author ]
    in http://example.org/books return
  <author>
  { $author }
}
```

```
        <name>{$author/bn}</name>
      </author>
    </book>
  }
</books>
```

Here the association is not queried by a path expression, but rather by a constraint in AsTMa! syntax.

Project page:

<http://astma.it.bond.edu.au/querying.xsp>

Implementation(s):

as part of the Perl XTM module, available via CPAN

Online demonstration:

<http://astma.it.bond.edu.au/query/>

4.4.3 Toma: Querying Topic Maps inspired by SQL

Due to its wide-spread acceptance, designing a query language with syntax and features similar to SQL is a natural choice. A first proposal for a Topic Maps query language inspired by SQL can be found in [158], that allows to query topics, topic-types, and associations using a mixture of SQL syntax and path expressions. E.g., the following query selects all books (i.e., topics classified as `Writing`) together with their authors:

```
select topic[book], topic[author]
from   topic-type["Writing"].topic[book],
       topic[book].assoc[a].topic[author],
       assoc-type["author-for-book"].assoc[a]
```

[158] also points out that a close alignment of query language and representation formalism is desirable, e.g., to allow the user to author a query with (nearly) the same tools used for authoring the topic maps itself.

Developed at Space Application Services¹⁵ in a project for the European Space Agency, Toma [213] is a more elaborate proposal for a Topic Maps query language along the lines of SQL and path expressions known from object-oriented query languages. It provides access to all Topic Maps concepts, including an implementation of the special semantics of the subsumption hierarchy. Information about a topic such as topic ID, basename, or subject identifier are accessed using a `.` notation as in object-oriented languages. E.g., `$topic.bn = 'Julius Caesar'` compares the basename of all topics selected by `$topic` with the string "Julius Caesar". Associations can be traversed using `->`, however only the associations with special semantics such as the instance-of or superclass-subclass association can also be traversed transitively, i.e., without knowing the length of the path to be traversed a-priori. Indeed, for traversing the subsumption hierarchy Toma provides the notation `$start.super(1..*)` that selects all super-classes of the current one. Instead of `1..*` one can specify any interval or single number to indicate how many superclass-subclass associations shall be traversed. A similar notation is available for instance-of associations.

A Toma expression implementing Query 1 is shown in the following:

¹⁵ <http://www.spaceapplications.com/>

```
select $book, $author, $author.bn
where $book.type(1..*).id = 'Writing'
      and author-for-book%->Writing = $book
      and author-for-book%->author = $author
```

This query selects all topics classified as `Writing` or one of its subtypes together with their authors and the basenames of the authors. The link between the book and the author is established by the `author-for-book` associating: a topic x is the author of another topic y , if x occurs in role `author` and y in role `Writing` in the same association (the use of `%a` in the query ensures that the same association is used).

Using the above mentioned constructs for querying the type hierarchy, Query 3 can be implemented easily: The following Toma query selects all topics that are neither used as type of another topic (i.e., part of the “ontology”) nor typed as `Translator`:

```
select $topic
where $topic.type(1..*).si.sir != 'http://example.org/books#Translator'
      and not exists ($t.type(1) = $topic)
      and not exists ($t.type(1..*) = $x and $topic.super(1..*) = $x)
```

In this query all topics are selected that neither (a) have the subject identifier `http://example.org/books#Translator` nor (b) are the type of some topic nor (c) are a sub-class of topic that is the type of some topic.

Project page:

<http://www.spaceapplications.com/toma/>

Implementation(s):

implementation not freely available

Online demonstration:

none

4.4.4 Path-based Access to Topic Maps

As for RDF and XML query languages, following the success of XPath great interest in path-based query languages for Topic Maps has been triggered (cf. [18] for an overview of such languages and a plea for inclusion of path navigation in the upcoming ISO Topic Maps query language). Actually, most of the current proposals for Topic Maps query languages use some form of path expressions with the notable exception of `tolog` (see Section 4.4.1). In the following, two languages focusing on providing navigational path expressions for accessing Topic Maps languages are investigated. Further proposals for path expressions as basis for querying Topic Maps are discussed in Section 4.4.2 and 4.4.3.

4.4.4.1 XTMPath

XTMPath [20] is an approach to define an easy-to-use language for accessing data stored in topic maps using XPath-like path expressions. The language is defined in [20] as part of the Perl XTM toolkit, [130] presents an easy overview. The core principle of the language is to use the way Topic Maps are serialized as XML documents in XTM [212] as orientation on how

to address parts of a topic map. E.g., the expression `topic[instanceOf/topicRef/@href = "#Historical_Novel"]` finds all topics that are (directly) typed as `Historical_Novel`. The path expression results from the way this information is represented in XTM:

```
<topic id="b1">
  <instanceOf>
    <topicRef xlink:href="#Historical_Novel"/>
  </instanceOf>
</topic>
```

So, given an XTM serialization of a topic maps finding the correct paths for addressing a part of that topic map is fairly easy.

However, one has to be aware of certain peculiarities of XTMPPath:

- Only a very small subset of the language constructs provided in XPath is currently supported by XTMPPath, mostly abbreviated syntax for child and descendant axis and some simple predicates.
- [20] argues that, in contrast to XPath, the XTMPPath processor essentially operates on data conforming to a single DTD (viz., the XTM DTD). This observation leads to treating the child axis in most cases as equivalent to the descendant axis. Only in some rare cases (e.g., for `instanceOf`) a difference between child and descendant is made.

Project page:

<http://cpan.uwinnipeg.ca/htdocs/XTM/XTM/Path.html>

Implementation(s):

as part of the XTM toolkit available from CPAN

Online demonstration:

none

4.4.4.2 TMPPath

TMPPath [42, 43] is an “experimental” language for accessing parts of a topic maps in a style similar to XPath. It has been developed by Dmitry Bogachev, partly as input to the ISO Topic Maps Query Language working group.

Just as XPath, it is designed to be used as embedded language that implements the selection of (or access to) parts of the topic maps, that can be used by the host language for further processing.

The latest version of TMPPath [43] provides a large number of constructs for this task. In contrast to the strict syntax of compound steps in XPath consisting in axis and node-test with an optional predicate, TMPPath mixes different styles of steps. E.g., the following TMPPath expression returns the basename of all authors, i.e., all topics that occur in the `author` role in a `author-for-book` association:

```
/topic[*;roleOf::author[is-author-of]/role::Writing]/bn::*[1]
```

The syntax `*;...` is one of the many shortcuts in TMPPath specifying a type condition before the semicolon and arbitrary further conditions after it.

TMPPath also provides means for binding variables using `for` clauses familiar from XPath. This allows, e.g., to return a list of all Writings together with their authors:

```
for $book in /topic[subjectIdentifier = "http://example.org/books#Writing"]
  for $author in /topic[*;roleOf::author[is-author-of]/role::Writing = $book]
    return list{$book/bn::*[1],$author/bn::*[1]}
```

Although obviously inspired by XPath, the lack of strict rules for representing the various TMPath step and predicate expressions leads to a rather complicated and hard to read syntax.

Project page:

<http://homepage.mac.com/dmitryv/TopicMaps/TMPath/TMPathRevisited.html>

Implementation(s):

unclear, not freely available

Online demonstration:

none

4.5 OWL Query Languages

4.5.1 OWL-QL

OWL Query Language (OWL-QL) combines a formal language for querying OWL ontologies with a protocol designed to support a dialogue between a querying agent and an answering agent [103]. It is based on the earlier DAML Query Language (DQL) developed by the Joint US/EU ad hoc Agent Markup Language Committee [102]. Although based on OWL and RDF, the language is quite generic and could easily be adapted to other logic based knowledge representation formalisms such as SWRL [137] or SCL¹⁶.

Query Language In OWL-QL, the query language itself is based on the standard notion of language statements (in this case OWL statements) in which some terms are replaced by variables [83]. Queries will often resemble standard conjunctive queries with the predicates being OWL classes and properties, and constants being OWL individuals [139]. An answer to a query consists of a *binding* for some or all of the variables in the query, i.e., a set of individual names from the ontology(ies) which, when substituted for the corresponding variables, give a set of statements that are entailed by the ontology(ies) being queried. E.g., a query of the form

$$\text{Person}(?x) \text{ worksFor}(?x, \text{W3C})$$

asks for persons who work for W3C, and if the ontology contains the statements $\text{Person}(\text{Jane})$ and $\text{worksFor}(\text{Jane}, \text{W3C})$, then binding Jane to $?x$ answers the query as under this binding the statements in the query are entailed by the ontology.

As usual, there are different kinds of variable: *must-bind* variables are those which must be bound to some individual name (in conjunctive queries, these are usually called distinguished variables); *don't-bind* variables are those for which no explicit binding is required (in conjunctive queries, these are usually called non-distinguished variables); and *may-bind* variables are those which may optionally bound to an individual name. The addition of may-bind variables doesn't increase the expressive power of the language (the same result could be

¹⁶ Common Logic Standard, <http://c1.tamu.edu>.

achieved by combining the results of queries using only must-bind and don't-bind variables), but may be convenient in some applications. Variables that are not bound to individual names in a query answer are treated as existentially quantified. For full details of the semantics of OWL-QL, the reader is again referred to [103].

In OWL-QL, standard taxonomic queries, e.g., retrieving all the super-classes of a given class, can be answered by using RDF properties in query atoms. For example, the query

```
subClassOf(Person, ?x)
```

would return all the super-classes of Person.

Given that the semantics of OWL-QL are based on entailment, answering OWL-QL requires, in general, the use of a theorem prover. One possible technique is to transform both query and ontologies into First Order Logic and use a FOL order theorem prover; details of an implementation based on this idea can be found at the DQL project for the Stanford Knowledge Systems Laboratory (<http://ks1.stanford.edu/projects/dql/>). Another technique is to reduce the conjunctive queries to standard retrieval queries that can be answered using a Description Logic (DL) reasoner [139, 249, 136]; details of an implementation based on this technique can be found at [117]. Details of a Jess based implementation can be found at [235].

Query Answering Protocol The OWL-QL query answering protocol is designed to cope with the wide variety of situations that might arise in a heterogeneous web environment:

- there may be many different kinds of *server* (a query answering service) with access to different kinds of information represented in many different formats;
- servers may have only partial information and may have limitations with respect to their performance (speed and/or completeness), the language they can handle (e.g., OWL Lite/DL/Full) and the kinds of query they are able to answer;
- the querying agent may or may not want or be able to specify all of the ontologies that should be used in answering a query;
- the querying agent may need only one answer, all possible answers, or something in between;
- the querying agent may consider some parts of the answer to be more important than others;
- and querying agents and servers may use different forms of surface syntax, e.g., RDF or XML.

In order to address these requirements, a query answer can be returned in one or more *bundles*, the query agent can specify an upper limit on the size of answer bundles, and a server can indicate the characteristics of its answer with respect to completeness and duplication. The query agent can also specify zero or more ontologies and (optionally) additional OWL statements that are to be used in computing the query answer. Moreover, the query language is specified using an abstract syntax for which many different serialisations are possible (e.g., an XML serialisation has been defined at the DQL project for the Stanford Knowledge Systems Laboratory, <http://ks1.stanford.edu/projects/dql/>).

When the query agent specifies one or more ontology, then only these ontologies can be used to compute answers.¹⁷ Alternatively, the query agent may specify a variable instead of an ontology. In this case the server is free to use any ontology it chooses to answer the query (the idea here is that the server can use arbitrary web accessible resources in order to find answers to such queries); the actual ontology used to answer the query may be returned as a binding for the variable, depending on whether it is a must-bind or may-bind variable.

If the query agent specifies an upper limit on the size of the answer bundle, then the size of an answer bundle returned by the server may range from zero up to this limit. An answer bundle from the server also includes either a process handle or a termination token. In case a process handle is returned, this can be used by a query agent (not necessarily the same agent) to continue the dialogue by requesting more answers to the query; a query agent can also terminate the dialogue at this point by sending the server the process handle with a termination request. A dialogue can be ended by the server using one of three different types of termination tokens: *end* simply indicates that no more answers will be provided; *none* explicitly asserts that all possible answers have been returned (i.e., the union of the answer sets in this and any preceding bundles constitute a complete answer to the query); and “rejected” indicates that the server is unable to answer the query. This last case covers a range of possibilities, including queries being outside the scope of a particular server (e.g., an OWL DL query sent to an OWL Lite server), or simply ill formed.

The specification of OWL-QL envisages servers with different kinds of behaviour regarding the generation of duplicate answers (although it does not specify a language mechanism whereby this information could be communicated to a query agent). A *non-repeating* server is one which guarantees not to duplicate answers during the course of a dialogue; a *terse* server is one that will not return redundant answers, where an answer is considered redundant if it is less specific than another answer (i.e., it has the same bindings for must-bind variables and a subset of the bindings for may-bind variables); a *serially-terse* server is one that will not return answers that are redundant with respect to already returned answers (but answers could be rendered redundant by a subsequent answer).

Project page:

<http://ks1.stanford.edu/projects/owl-ql/>

Implementation(s):

available from the project page

Online demonstration:

<http://onto.stanford.edu:8080/>

¹⁷ Note that these ontologies may import others using OWL’s import mechanism.

Chapter 5

Evaluation Results

In Appendix B, the detailed evaluation of the query languages discussed above is shown. Here, some of the most striking results are highlighted.

The first and most obvious observation that can be derived from the discussion of the query languages in Chapter 4 and the feature evaluation from Appendix B is the great **variety** of proposed Web query languages ranging from path languages providing only the most basic means for data access (XTMPPath, Section 4.4.4.1; RDFPath, Section 4.3.2.5) over similarly basic languages for extracting multiple data items at once (RDQL, Section 4.3.1.3; XQL, Section 4.2.1.3), languages with and without ontology support to computationally complete languages with general recursion that are able to query data in any of the representation formalisms considered here (Xcerpt, Section 4.2.2.6; XQuery, Section 4.2.1.5; TRIPLE, Section 4.3.5.8) or constrained natural language for querying Semantic Web data (Metalog, Section 4.3.4.1). This variety is also reflected in the ability of the different query languages to infer data: some languages do not consider inference, most provide at least some restricted form of inference, e.g., means for computing the transitive closure of a relation, some allow for a Datalog-like specification of intensional data.

Also rather obvious from the results, although not commonly acknowledge is that similar approaches occur for XML, RDF, and Topic Maps querying, e.g., basic path-based languages are proposed for all cases and query languages inspired by logic programming languages such as Prolog or Datalog are often the ones providing the richest expressiveness in each area. This is particularly true when considering the reasoning abilities of the surveyed languages: in all areas languages with no actual reasoning at all, with very limited reasoning abilities for implementing specificities of the underlying representation formalism, and with strong reasoning support, e.g., by means of general recursion and Horn logic clauses, can be found.

This points to a common classification scheme for the query languages surveyed so far using the reasoning abilities of a language as distinguishing property. Such a classification scheme is proposed and discussed in the following section.

5.1 A Classification Scheme for Web Query Languages

In the tradition of the seminal papers by Codd [84] and later by Chandra and Harel [71, 72], query languages for (relational) data base have often been characterized by their expressive-

ness (or completeness under the relational algebra) and evaluation complexity. However, many recent proposals for Web query languages have acknowledge previous results (e.g., in [6]) suggesting that such classification schemes have to be altered for a Web context by providing computationally complete languages (e.g., XQuery [41], XSLT 2.0 [150], Xcerpt [230], XPathLog [187], TRIPLE [237], tolog [112]).

For Edutella [195], an RDF-based peer-to-peer (P2P) infrastructure, a language for exchanging queries, dubbed RDF-QEL, among the nodes in the P2P system has been developed. This language is based on Datalog but to support a wide range of devices and implementations with differing capabilities, five language layers are proposed distinguished by increasing complexity:

- *Rule-less queries* are queries without rules (or equivalent constructs).
- *Conjunctive queries* allow only a single, non-recursive rule per predicate.
- *Disjunctive queries* can use several rules for defining a predicate but may not be recursive in any way.
- *Linear recursive queries* may contain predicates defined by linear recursion.
- *General recursive queries* may contain arbitrary Datalog predicates.

This classification scheme is very useful to estimate the processing capabilities required for evaluating a query. However, considering the results of this survey, as the traditional classifications based on query complexity, it proves not to deliver an interesting and revealing classification of the languages. This can be attributed to the fact that the RDF-QEL classification has been defined for queries not for languages. E.g., the class of linear recursive queries is certainly interesting, however among the languages surveyed here there is no language that supports only linear recursion.

Therefore, a novel classification scheme for Web query languages is proposed here. The core classification feature used in this scheme is the Semantic Web “fitness”, the reasoning abilities of a language. Four classes of languages are proposed, the third one divided into two subclasses, yielding the hierarchy depicted in Figure 5.1:

Class 1: Selection-Only Languages. The main characteristic of these languages, a typical example of which is XPath [78], is the lack of construction abilities, i.e., they are only able to specify which part of the input to be *selected* by the query. A direct consequence of this restriction is the missing of means for construction of intermediate results such as views, rules or functions. Also most of these languages operate on a single document (or similar collection of data given by the context).

Although these restrictions limit the expressiveness severely, they also allow for an easy implementation and, at least in some cases, efficient evaluation of the languages. Furthermore, these languages are often used as part of other technologies or even other query languages, hence allowing the use of the same basic access functionality and syntax over a wide range of technologies.

Class 2: Non-recursive Languages. In contrast to the selection-only languages, these class of languages provides some means of *construction*, often realized by nested queries. However, recursion as needed for inference queries and the traversal of arbitrary-length paths in a structure are either missing entirely or only available on some predefined relations

(e.g., parent-child relation in XML, but no traversal of arbitrary-length ID/IDREF or XLink relations).

Also, these languages do not specifically support the extended semantics provided by ontology languages such as RDFS or OWL, although in some cases (e.g., RDQL) certain implementations provide very limited ontology support as part of the storage model.

Such languages can express not only selection, extraction, restructuring, and often also reduction queries, but also some inference and combination queries, where a fixed upper bound on the size of the inferred data exists.

Typical examples of this class are, e.g., RDQL [232] and XML-QL [94].

Class 3a: Ontology-aware non-recursive languages. These languages specifically support the use of ontology information for querying, but do not allow the use of general recursion. A well-known example for such a language is RQL [148], which employs ontologies for typing and querying but limits the traversal of arbitrary-length paths to the subsumption hierarchy defined in the ontology.

Class 3b: Recursive languages without specific ontology support. A large number of query languages, in particular for querying XML (e.g., XSLT, XQuery), falls into this class: they provide the ability to express recursive queries on top of the capabilities of Class 2 languages. However, no specific support for ontology reasoning is given. This is not so much a limitation on the expressiveness of the query languages (most of the languages in this class are computationally complete anyway), as on the convenience (and potentially efficiency) for expressing queries relying on ontology information. The mechanisms for inferring knowledge from an ontology describing the queried data have to be explicitly stated in a query.

Often such languages are the basis of extensions (in form of libraries or true language extensions) for supporting ontology reasoning, i.e., the basis of Class 4 languages.

Class 4: Ontology-aware recursive languages. Only languages that support both general recursion (or equivalent operations) and the specificities of some ontology languages such as RDFS or OWL are included in this class. Representative languages are, e.g., Xcerpt, TRIPLE, or tolog.

This comparison scheme is obviously tailored to querying the kind of data envisioned to be predominant in a Semantic Web setting: Heterogeneous, highly, but often inconsistently structured. Flexible means for programmatic manipulation of such data are called for. This entails, in particular, the ability to query and traverse arbitrary-length paths of related items in the data, both if the relation is expressed in the structure (e.g., parent-child relationship in XML) or established by other means (e.g., ID/IDREF, XLink, based on foreign keys such as URIs).

Traversing arbitrary-length paths of related items is one of the most basic inference query that is required by many use cases proposed for the Semantic vision. Also central to the idea of the Semantic Web is the use of formally defined vocabularies that allow a more precise automated “understanding” of the data described.

The classification scheme proposed in this section uses these two observations for providing an structured view on the languages considered in this survey.

Figure 5.1 Excerpt of the Surveyed Languages in Classification Scheme

Ontologies & Full Recursion	Class 4 <ul style="list-style-type: none"> ontology support predefined constructs for relations with special semantics closure on arbitrary relations <small>.Syntactic Web? Xcerpt for SW</small>	Class 4 <ul style="list-style-type: none"> ontology support full recursion closure on arbitrary relations <small>TRIPLE (RDFS only)</small>	Class 4 <ul style="list-style-type: none"> ontology support full recursion closure on arbitrary associations 	
Limited Recursion	Class 3 <ul style="list-style-type: none"> full recursion allows querying relations with special semantics allows closure on relations <small>XSLT 2.0 XQuery 1.0 Xcerpt ...</small>	Class 3 <ul style="list-style-type: none"> no ontology support full recursion or closure <small>RDFQL Versa R-DEVICE</small>	Class 3 <ul style="list-style-type: none"> ontology support no general recursion <small>RQL SeRQL OWL-QL</small>	Class 3 <ul style="list-style-type: none"> no ontology support full recursion or closure ontology support no general recursion <small>tolog AstMe? Toma</small>
Querying Semantic Web	Class 2 <ul style="list-style-type: none"> multiple documents nested queries or limited views but no recursion querying of different serializations of RDF/TopicMaps possible <small>XML-QL Lorel</small>	Class 2 <ul style="list-style-type: none"> (mostly) multiple documents no recursion querying an RDF model <small>SquishQL RDQL RDFPath</small>	Class 2 <ul style="list-style-type: none"> (mostly) multiple documents no recursion querying on TopicMap model 	
Addressing	Class 1 <ul style="list-style-type: none"> limited or no views only limited selection queries single-document queries specific to one serialization of RDF/TopicMaps <small>XSLT 1.0 XPath 1.0 XQL Lorel</small>	Class 1	Class 1 <ul style="list-style-type: none"> no views, no recursion only limited selection queries single-document <small>XTM: :Path</small>	
	XML	RDF	TopicMaps	

5.2 Observations on the State of the Art of Web Query Languages

Aside of the classification proposed in the previous section, there are a few additional observations on the status of Web query languages that can be derived from the above comparison:

- Varying Maturity Level:** The query languages surveyed in this report vary noticeably in the level of maturity. As to be expected, query languages for RDF and Topic Maps are in general less evolved than XML query languages that have been investigated in academia and industry for several years. But also, e.g., the proposed RDF query languages differ quite noticeably in their maturity level, some still barely more than quickly drafted proposals, some already in production use.
- Intense and Early Standardization Activity:** In all three areas, but particularly for XML and Topic Maps query languages standardization activity from various organizations such as W3C or ISO precedes or runs parallel to early implementations and research activity. This can lead to a premature focus on alignment with use cases and requirements proposed by the standardization bodies.
- Layering of Query Languages:** A possible explanation in the variety observed in the query languages proposals surveyed here is that some of the query languages are limited to a specific task such as the selection of data in a XML, RDF, or Topic Maps structure. Often such limited query languages are then used as a separate access layer in full transformation and query languages. The typical example for such layering is the use of XPath for accessing nodes in an XML document in XSLT, which provides more advanced restructuring and transformation capabilities.
- Approximation and Reasoning as an Emerging Issue:** At least two issues are receiving increasing interest recently, both motivated by the characteristics of data observed or expected in the (Semantic) Web:

(1) Not all data is structured properly, quite a lot of interesting information can only be deduced from full-text processing, and there is (not yet) a common understanding of how to structure data properly. This leads to the desire for features in the style of information retrieval systems that (a) allow the processing of full-text data included in the structure and (b) can be extended to allow queries where the structure of the data is only approximately known (e.g., whether a data item is represented as element or attribute).

(2) Combining such issues with formal vocabularies (from simple thesauri to ontologies described, e.g., in OWL full) requires some ability to reason about the provided data, e.g., for discovering relations between data items not explicitly represented in the structure.

Finally, one should note the similarity of common issues found to be interesting for query languages from all three areas. E.g., in all cases, one has to consider traversing arbitrary-length paths in the relations provided by the representation formalism. Also in all cases, ontology information and similar reasoning techniques can increase the recall of a query in face of heterogeneous descriptions of the data.

Chapter 6

Conclusion and Outlook

This survey presents a unique look on query languages based upon the diverse formats for data representation already used or expected to be used in the (Semantic) Web. It is illustrated that there is a large number of issues common to query languages in each of these areas that goes well beyond general design considerations of query languages but is rooted in the characteristics of the Web context. The most prominent issues are the ability to handle heterogeneous data both in structure and content, to support description of the same or similar information types using different vocabularies, to allow incomplete or approximate specifications of queries, and the consideration of reasoning abilities to be able to integrate, mediate, and enrich the data provided.

The results of the evaluation show that a unified classification scheme for XML, RDF, and Topic Maps can be derived that is both meaningful and interesting for understanding the different proposals and their intended usage scenarios. Such a classification scheme together with the detailed results presented in this report can help identifying interesting languages for varying requirements and provides a better insight in the state of the art of Web query languages.

In the perception of the authors, these results stress the need for a query language that is able to handle all these representation formalisms and the plethora of serialization formats proposed for them.

To understand further the requirements for such a language, the REVERSE I4 working group is investigating design principles (a preliminary report can be found in [60]) and use cases for such a query language. Xcerpt [230] represents a first proposal for a language targeted at the flexibility and reasoning capabilities required in this setting.

Acknowledgements. We would like to thank Wolfgang May for reviewing a draft of this deliverable and providing numerous invaluable comments on how to improve both its presentation and content.

This research has been funded by the European Commission and by the Swiss Federal Office for Education and Science within the 6th Framework Programme project REVERSE number 506779 (cf. <http://reverse.net>).

Appendix A

A Brief History of RDF Serialization Formats

A.1 Introduction

The *Resource Description Framework* (RDF) is a language for making simple statements about resources on the World Wide Web in form of a graph of nodes and edges representing the resources, their properties and values. A standard XML syntax for serializing RDF graphs, RDF/XML, exists. Yet, in the five years after the initial release of RDF, numerous alternative serialization formats have been proposed. This report attempts to present an overview of the different proposals and the motivations behind their creation.

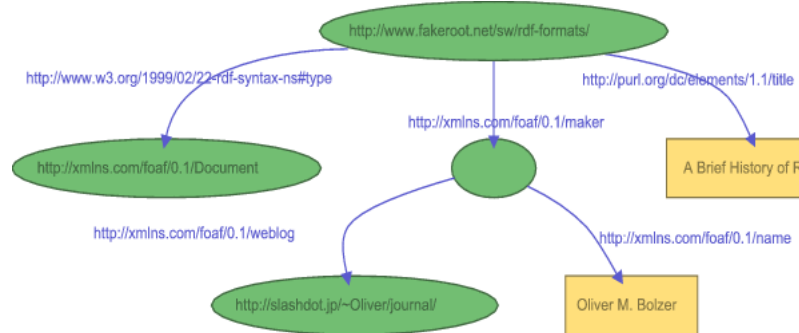
The serialization formats presented here can be categorized into three general classes, depending on the use of XML and how a RDF graph is seen: either as fully connected graph or as collection of subject-property-object triples. RDF/XML and its simplification attempts try to map the nodes and edges of a RDF graph directly to elements in a XML tree. The plain-text formats deriving from N3 concentrate on the individual triples that make up the graph and record them in a non-XML form. The newest generation of serialization formats, TriX and RXR also focus on the triples, but use XML to specify them. A fourth class of formats, specially designed for embedding RDF data into XHTML and other XML languages are not considered in this report.

The RDF graph presented in Figure A.1 will be used throughout this report, serialized into each format. It describes this report and its author using the Friend-Of-A-Friend [55] and Dublin Core [143] vocabularies.

A.2 RDF/XML: The W3C Recommendation

When RDF became a W3C Recommendation for the first time in 1999, [167] defined together with the formal RDF model a XML based syntax for serializing RDF graphs: RDF/XML. Because of the differences in the underlying information models of RDF and XML, one being an edge-and-node-labeled directed graph of resources and properties identified by URIs and the other being a node-labeled tree of elements and attributes identified by the combination of names-

Figure A.1 A sample RDF Graph



pace and tag name, the specification proposed a mapping where both resources and properties were converted to XML elements and nested into each other. By using the XML namespace mechanism, it was possible to split an URI into two parts and form an XML element name. The main concept behind RDF/XML later became to be known as *striping* [53], as resources and properties alternate in the nested XML structure.

Query 20. RDF/XML describing this report and it's author.

```
<?xml version="1.0"?>
<rdf:RDF xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
  xmlns:dc="http://purl.org/dc/elements/1.1/"
  xmlns:foaf="http://xmlns.com/foaf/0.1/" >
  <rdf:Description rdf:about="http://www.fakeroot.net/sw/rdf-formats/">
    <rdf:type>
      <rdf:Description rdf:about="http://xmlns.com/foaf/0.1/Document" />
    </rdf:type>
    <dc:title>A Brief History of RDF Serialization Formats</dc:title>
    <foaf:maker>
      <rdf:Description>
        <foaf:name>Oliver M. Bolzer</foaf:name>
        <foaf:weblog>
          <rdf:Description rdf:about="http://slashdot.jp/~Oliver/journal/" />
        </foaf:weblog>
      </rdf:Description>
    </foaf:maker>
  </rdf:Description>
</rdf:RDF>
```

Example 20 is a RDF/XML document for the RDF graph from Figure A.1. `rdf:Description` elements represents the resource identified by the URI in it's `rdf:about` attribute. Blank nodes do not have a `rdf:about` attribute, but can optionally have a `rdf:nodeID` attribute to distinguish them from other blank nodes in the same graph. Direct children of a `rdf:Description` elements are the properties describing the resource. For these, the URI identifying the property is used as the element name, by splitting the URI into a prefix and a suffix, used as XML namespace and local part.

RDF/XML allows to shorten the serialization using various abbreviations. For instance, a property with a literal as object can be expressed using a XML attribute. Also the resource that is the object of an statement can be named in the `rdf:resource` attribute of the element for

the property, instead of an `rdf:Description` child element. Also the value of a resource's `rdf:type` property can be used as the resource's element name instead of `rdf:Description`. Example 21 is also a RDF/XML document for Figure A.1, but much shorter than the previous example, through the use of abbreviations.

Query 21. RDF/XML with abbreviations.

```
<?xml version="1.0"?>
<rdf:RDF xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
  xmlns:dc="http://purl.org/dc/elements/1.1/"
  xmlns:foaf="http://xmlns.com/foaf/0.1/" >
  <foaf:Document rdf:about="http://www.fakeroot.net/sw/rdf-formats/"
    dc:title="A Brief History of RDF Serialization Formats" >
    <foaf:maker>
      <rdf:Description foaf:name="Oliver M. Bolzer">
        <foaf:weblog rdf:resource="http://slashdot.jp/~0liver/journal/" />
      </rdf:Description>
    </foaf:maker>
  </rdf:Description>
</rdf:RDF>
```

Many more possibilities to serialize the same graph using RDF/XML exist, because the aggregation of multiple statements about a single resource into children of a single `rdf:Description` element are not mandatory. Also, instead of deeply nesting resources and properties, it is possible to create a shallow structure with all `rdf:Description` elements directly under the root. Example 22 is another serialization of Figure A.1, this time in a very verbose way.

Query 22. Very verbose RDF/XML.

```
<?xml version="1.0"?>
<rdf:RDF xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
  xmlns:dc="http://purl.org/dc/elements/1.1/"
  xmlns:foaf="http://xmlns.com/foaf/0.1/" >
  <rdf:Description rdf:about="http://www.fakeroot.net/sw/rdf-formats/">
    <dc:title>A Brief History of RDF Serialization Formats</dc:title>
  </rdf:Description>
  <foaf:Document rdf:about="http://www.fakeroot.net/sw/rdf-formats/">
    <foaf:maker rdf:nodeID="oliver">
  </rdf:Description>
  <rdf:Description rdf:nodeID="oliver">
    <foaf:name>Oliver M. Bolzer</foaf:name>
  </rdf:Description>
  <rdf:Description rdf:nodeID="oliver">
    <foaf:weblog rdf:resource="http://slashdot.jp/~0liver/journal/" />
  </rdf:Description>
</rdf:RDF>
```

Examples 20-22 all represent the exact same RDF graph. After being processed by a RDF/XML parser, there should be no difference between them. But from the view point of XML tools, each one is a completely different XML document. Because of this syntactic variability, it is very difficult to use standard XML tools like XPath, XSLT and XQuery with arbitrary RDF/XML documents.

Furthermore, several other problems have been identified with RDF/XML, which all led to the development of the subsequent serialization formats. The following are among the most notable:

- It is impossible to distinguish an XML element for a RDF node from one for a property without knowledge about the striping.
- The triples that make up the RDF graph are hard to make out.
- Many different things are used to to represent URIs: element names, attribute names and attribute values.
- It is impossible to specify a single DTD or XML Schema that validates all RDF/XML documents.
- Difficult to read for humans.

In 2001, the RDF Core Working Group was created, partly to fix the RDF/XML syntax and clean up the whole specification. This effort led to the revised RDF/XML Syntax Specification [27], which became a W3C Recommendation in early 2004. However, though the specification has undergone a major clean up and the syntax is now specified in a cleaner and much more concise manner, the basics have not changed and the problems arising from RDF/XML's structure have not been solved. [25] gives an overview of the problems of the first RDF/XML specification as well as the issues identified during the revision process.

A.3 Simplified Syntaxes for RDF/XML

A.3.1 Unstriped Syntax

Only few month after the publication of [167], Tim Berners-Lee started experimenting with simplifications of RDF/XML. In [35] he considered a modification of RDF/XML without the node/property striping, named "Unstriped Syntax". In it, XML elements are only used for the edges in the RDF graph, with the subjects specified using the newly introduced `rdf:for` attribute.

Query 23. Single Statement using the Unstriped RDF/XML Syntax.

```
<dc:title rdf:for="http://www.fakeroot.net/sw/rdf-formats/">A Brief History of RDF
Serialization Formats</dc:title>
```

Alternately, a default subject for all nested elements can be given using a `rdf:about` attribute on the parent element, similar to RDF/XML. Blank nodes and deep nesting of elements are handled as in RDF/XML.

Query 24. Figure A.1 serialized using the Unstriped Syntax.

```
<someelement rdf:about="http://www.fakeroot.net/sw/rdf-formats/">
  <rdf:type rdf:about="http://xmlns.com/foaf/0.1/Document" />
  <dc:title>A Brief History of RDF Serialization Formats</dc:title>
  <foaf:maker>
    <foaf:weblog rdf:about="http://slashdot.jp/~oliver/journal/" />
    <foaf:name>Oliver M. Bolzer</foaf:name>
  </foaf:maker>
</someelement>
```

However, lacking a suitable parent element (somelement in the above example), use of the `rdf:Description` element is recommended, undermining the Unstripped Syntax's basic principle that elements are only to be used for properties.

Because it addresses only the striping issue and none of the other problems with RDF/XML, the Unstripped Syntax has not been pursued further than it's "strawman draft" status.

A.3.2 Simplified Syntax

Inspired by the Unstripped Syntax, Sergey Melnik followed up with another simplified RDF/XML syntax [191], in which each element's type, whether it is a node or an edge of the RDF graph, can be determined by looking at it's attributes. As with the Unstripped Syntax, each XML element by default denotes an edge of the graph. The subject is specified in a `rdf:for` attribute and it's object in a `rdf:resource` attribute or in case of a blank node, using child elements. In absence of a `rdf:for` attribute, the subject is defined to be the parent element's object, instead of an explicitly set "default subject", removing the need to use an extra element to denote a resource.

Query 25. Figure A.1 serialized using the Simplified Syntax.

```
<rdf:type rdf:for="http://www.fakeroot.net/sw/rdf-formats/" rdf:resource="http://xmlns.com/foaf/0.1/Document" />
<dc:title rdf:for="http://www.fakeroot.net/sw/rdf-formats/">A Brief History of RDF Serialization Formats</dc:title>
<foaf:makes rdf:for="http://www.fakeroot.net/sw/rdf-formats/">
  <foaf:name>Oliver M. Bolzer</foaf:name>
  <foaf:weblog rdf:resource="http://slashdot.jp/~Oliver/journal/" />
</foaf:makes>
```

Only when an element has a `rdf:instance` attribute, does it denote a node identified by the URI in the attribute's value. In such a case, the element's name is taken to be the node's `rdf:type`. Though this feature reintroduces striping, it is explicit and easily detectable. Example 26 shows parts of our example RDF graph utilizing this feature.

Query 26. Simplified Syntax using `rdf:instance`.

```
<foaf:Document rdf:instance="http://www.fakeroot.net/sw/rdf-formats/">
<dc:title>A Brief History of RDF Serialization Formats</dc:title>
<foaf:makes>...</foaf:makes>
</foaf:Document>
```

A.3.3 XMP

Faced with the need to embed meta-data into the various media formats produced by it's tools, Adobe decided to adopt RDF as the core of it's "Extensible Metadata Platform" [7]. Instead of going for the full RDF/XML format, Adobe uses only a proper subset, disallowing XML literals and reification in order to simplify processing and reduce the complexity of the expressible RDF graphs. As these removed features are rarely used, it is likely that the majority of RDF/XML documents on the Internet are also valid XMP data.

A.3.4 Normalized RDF

In 2001, Jonathan Robie demonstrated that a "normalized" form of RDF/XML could be effectively queried using XQuery, an XML query language without any knowledge about RDF. In [224]

he argued, that by standardizing on one of the many possible syntactic variants of RDF/XML, it would be possible to use standard XML tools to effectively query, transform and otherwise process RDF/XML documents.

Going through several refining steps in his paper, Robie arrived at a flat form, where all statements about a single resource are grouped together and all groups put under a common parent, thus avoiding deep nesting of statements.

Query 27. Figure A.1 in “normalized” RDF/XML.

```
<rdf:Description about="http://www.fakeroot.net/sw/rdf-formats/">
  <rdf:type>http://www.w3.org/1999/02/22-rdf-syntax-ns#type</rdf:type>
  <dc:title>A Brief History of RDF Serialization Formats</dc:title>
  <foaf:maker>http://fakeroot.net/staff/Oliver</foaf:maker>
</rdf:Description>
<rdf:Description about="http://fakeroot.net/staff/Oliver">
  <foaf:name>Oliver M. Bolzer</foaf:name>
  <foaf:weblog>http://slashdot.jp/~Oliver/journal/</foaf:weblog>
</rdf:Description>
```

Being only a technical demonstration for the possibility of querying RDF using XQuery, details like the difference between resources and literals as objects or blank nodes are not addressed. In Example 27 a new URI had to be assigned to identify the maker of the document.

A.3.5 RxML

RxML [239] is a serialization format by Adam Souzis created as component of his Rx4RDF suite of RDF-related technologies. It is unique in its consistent use of XML element names to encode all URIs extending the way properties are handled in RDF/XML to subjects and objects. Each child of the root `rx:rx` element specifies a resource describe in the RxML document. Its children are in turn the properties and the grandchildren are objects: either text children for literals or empty elements for resources. Nesting is not allowed, limiting the maximum depth of the XML tree to 3. Blank nodes are identified by resources whose URIs begin with ‘`bnode:`’.

Query 28. Figure A.1 in RxML

```
<rx:rx xmlns:rx='http://rx4rdf.sf.net/ns/rxml#'
  xmlns:bnode='bnode:'
  xmlns:rdf='http://www.w3.org/1999/02/22-rdf-syntax-ns#'
  xmlns:fakeroot='http://www.fakeroot.net/sw/'
  xmlns:dc='http://purl.org/dc/elements/1.1/'
  xmlns:foaf='http://xmlns.com/foaf/0.1/' >
  <fakeroot:rdf-formats>
    <rdf:type><foaf:Document /></rdf:type>
    <dc:title>A Brief History of RDF Serialization Formats</dc:title>
    <foaf:maker><bnode:Oliver /></foaf:maker>
  </fakeroot:rdf-formats>
  <bnode:Oliver>
    <foaf:name>Oliver M. Bolzer</foaf:name>
    <foaf:weblog><journal xmlns="http://slashdot.jp/~Oliver/" /></foaf:weblog>
  </bnode:Oliver>
</rx:rx>
```

While the consistent use of XML element names for URIs seems an elegant solution, it is accompanied by a fatal problem: some URIs can't be expressed due to restrictions in the characters allowed for element names in XML. In Example 28, the document's URI (<http://www.fakeroot.net/sw/rdf-formats/>) can't be turned into a XML element name because of the trailing '/' that had to be omitted.

A.4 Plain-Text Formats

A.4.1 Notation 3

Meanwhile, giving up on a usable XML syntax for RDF, Tim Berners-Lee proposed Notation 3, also known as N3 [32]. Contrary to previous serialization formats, N3 is not a XML based format. Born out of a pseudo-syntax people started using in various discussion forums instead of RDF/XML, N3 focuses on the triples that make up a RDF graph and writes them down in a straight manner: subject, property, object .

Query 29. A single Statement in N3.

```
<http://www.fakeroot.net/sw/rdf-formats/> <http://purl.org/dc/elements/1.1/title>
"A Brief History of RDF Serialization Formats" .
```

Two shortcuts are provided to combine several statements. A semicolon introduces another property about the same subject and a comma introduces another object with the same property and subject. Blank nodes are identified by using square brackets as objects, putting the statements about that blank node inside the brackets. Additionally, N3 allows the use of short prefixes to abbreviate long URIs, similar to the namespace prefixing mechanism in XML. Also the very common `rdf:type` predicate can be abbreviated to just an `a`. [33] gives an excellent introduction to the basics of N3.

Query 30. Figure A.1 in N3.

```
@prefix rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#> .
@prefix dc: <http://purl.org/dc/elements/1.1/> .
@prefix foaf: <http://xmlns.org/foaf/0.1/> .
<http://www.fakeroot.net/sw/rdf-formats/> a foaf:Document;
  dc:title "A Brief History of RDF Serialization ...";
  foaf:maker [
    foaf:weblog <http://slashdot.jp/~Oliver/journal>;
    foaf:name "Oliver M. Bolzer" ] .
```

N3 supports reification by quoting statements within curly braces. A quoted statement can then be used as subject for another statement.

Query 31. Reification in N3.

```
{ ex:Moby_Dick foaf:maker ex:Oliver .} ex:trustable ex:false .
```

N3 is also not just an serialization format for RDF graphs. It has additional support beyond RDF, allowing whole graphs to be quoted as well as formulation of rules and queries using variables and quantification. While most people only think of N3 as a serialization format, some people think of N3 as a rule language, while others consider it a query language. To

avoid confusion, attempts have been made to define subsets of N3 according to capability: N3 RDF, N3 Rules and full N3.

Being easy to read for both humans and machines, N3 was quickly adopted by the Semantic Web community as the format used for online discussions about RDF. Today, various tools and query language implementations for RDF accept N3 as input and output format together with RDF/XML.

A.4.2 N-Triples

N-Triples is a minimalist subset of N3, only allowing one triple per line without any abbreviations. Designed for RDF Test Cases [123], it is intended to be extremely easy to parse and generate by scripts. To avoiding any nesting, Blank nodes need to be identified by a temporary identifier starting with `_:`. N-Triples neither supports URI abbreviation, nor reification.

Query 32. Figure A.1 in N-Triples.

```
<http://www.fakeroot.net/sw/rdf-formats/> <http://www.w3.org/1999/02/22-rdf-syntax-ns#type>
<http://xmlns.org/foaf/0.1/Document> .
<http://www.fakeroot.net/sw/rdf-formats/> <http://purl.org/dc/elements/1.1/title>
"A Brief History of RDF Serialization Formats" .
<http://www.fakeroot.net/sw/rdf-formats/> <http://xmlns.org/foaf/0.1/maker> _:a .
_:a <http://xmlns.org/foaf/0.1/weblog> <http://slashdot.jp/~Oliver/journal> .
_:a <http://xmlns.org/foaf/0.1/name> "Oliver M. Bolzer" .
```

Because of its precise and simple syntax and straight accordance with the RDF's concept of triples ([155]), N-Triples is often encountered in introductory documents about RDF, such as [179].

A.4.3 Quads

When aggregating RDF statements from multiple sources and saving them locally, tracking the origin, or context, of each statement becomes more and more important. Some storage systems store the URI of origin together with each triple, forming a "quad". Quads [229] is an extension of N-Triples, adding an optional fourth element for such context information.

Query 33. A single statement that originated at <http://www.fakeroot.net/sw/SampleG.rdf>, in Quads

```
<http://www.fakeroot.net/sw/rdf-formats> <http://purl.org/dc/elements/1.1/title>
"A Brief History of RDF ..." <http://www.fakeroot.net/sw/SampleG.rdf> .
```

The specification of Quads includes further extensions such as "compound names" and statement terminators other than the dot, but the semantics of them are not further explained.

A.4.4 Turtle

While more and more RDF-related tools adopted N3 in addition to RDF/XML, most of them implemented only a ad-hoc subset of N3, leaving out some of the more complex features that go beyond the RDF model. In light of such development, Dave Beckett proposed in the end of 2003 a new plain-text format Turtle ([14],[26]), extending N-Triples with some of the commonly used

and well understood features of N3, while staying within the RDF model. Turtle deliberately skips support for reification.

Among the features taken from N3 are short prefixes for long URIs and the abbreviations using commas and semicolons, as well as blank node creation using square brackets and collections. Also the default character encoding has been changed from US-ASCII to UTF-8.

Query 34. Figure A.1 in Turtle.

```
@prefix rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#> .
@prefix dc: <http://purl.org/dc/elements/1.1/> .
@prefix foaf: <http://xmlns.org/foaf/0.1/> .
<http://www.fakeroot.net/sw/rdf-formats/> a foaf:Document;
dc:title "A Brief History of RDF Serialization ...";
foaf:maker [
  foaf:weblog <http://slashdot.jp/~0liver/journal>;
  foaf:name "Oliver M. Bolzer" ] .
```

Example 34 is identical to Example 30 given above for N3 above, due to Turtle being mostly a subset of N3.

A.4.5 TriG

TriG [39] is the newest in the plain-text line of formats descending from N3, proposed by Chris Bizer. Dubbed as “a compact and readable alternative to the XML-based TriX” [39], TriG extends Turtle beyond the RDF model by adding support for serializing multiple graphs in one file, with the ability to give each a distinct name [65].

Query 35. Figure A.1 with additional name in TriG.

```
@prefix rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#> .
@prefix dc: <http://purl.org/dc/elements/1.1/> .
@prefix foaf: <http://xmlns.org/foaf/0.1/> .
@prefix fakeroot: <http://www.fakeroot.net/sw/> .
fakeroot:SampleG { fakeroot:rdf-formats/ a foaf:Document;
                    dc:title "A Brief History of RDF Serialization ...";
                    foaf:maker [
                      foaf:weblog <http://slashdot.jp/~0liver/journal>;
                      foaf:name "Oliver M. Bolzer" ] .
                    }
}
```

Example 35 shows the graph from Figure A.1, given the name *http://www.fakeroot.net/sw/SampleG*. Its triples are grouped together using curly braces, with the name prepended.

A.5 Triple-based XML Formats

During 2003, while completing the revision of RDF/XML, Dave Beckett summarized the inherent problems of RDF/XML and collected requirements and ideas for new serializations formats in [25]. Based on the experiences with N3 and RPV (see below), Beckett argued, that such a format should be closely based on the RDF graph via the terminology in [155] and be minimal in the number of alternate forms for the same RDF graph.

Though not a format proposal, [25] defines the up-to-date most extensive list of requirements for a new XML format.

A.5.1 RPV

In 2002, despite the increasing popularity of N3, Tim Bray still felt the need for a XML based format and created RPV [49] in order to leverage XML's diverse assets, such as support for Internationalization and widely-deployed base of software. The goal was to create a format that was entirely unambiguous and highly human-readable, by emphasizing the triples that make up a RDF graph.

RPV takes a subject centric view on triples, collecting statements about a single resource together, similar to the Normalized RDF approach by Robie. Instead of using property names as element names, RPV utilizes only two tags: R (as in resource) and PV (as in property/value). URIs are specified using the attributes *r*, *p* and *v* for resource, property and value, respectively. Blank nodes are emulated by giving a R element an *id* attribute but no *r* attribute.

Query 36. Figure A.1 in RPV.

```
<R r="http://www.fakeroot.net/sw/rdf-formats/">
  <PV p="http://www.w3.org/1999/02/22-rdf-syntax-ns#type" v="http://xmlns.org/foaf/0.1/Document" />
  <PV p="http://purl.org/dc/elements/1.1/title">A Brief History of RDF ...</PV>
  <PV p="http://xmlns.org/foaf/0.1/maker" v="#oliver" />
</R>
<R id="oliver">
  <PV p="http://xmlns.org/foaf/0.1/weblog" v="http://slashdot.jp/~oliver/journal" />
  <PV p="http://xmlns.org/foaf/0.1/name">Oliver M. Bolzer</PV>
</R>
```

In order to abbreviate long URIs, RPV allows the use of the attributes *rBase*, *pBase*, and *vBase*, providing base URIs similar to `xml:base` for resource, property and value, respectively. These bases apply to the element and all contained elements. As only one such base is allowed for each type, the usability is very limited in situations using vocabularies with varying prefixes.

Query 37. RPV with abbreviated URIs.

```
<R r="http://www.fakeroot.net/sw/rdf-formats/"
  pBase="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
  vBase="http://xmlns.org/foaf/0.1/" >
  <PV p="type" v="Document" />
  <PV p="http://purl.org/dc/elements/1.1/title">A Brief History of RDF ...</PV>
  <PV p="http://xmlns.org/foaf/0.1/maker" v="#oliver" />
</R>
```

RPV supports reification using the *rpv* attribute on a R element, pointing to another R element's *id*. In Example 37, in the first R element, identified as *foo* by the *id* attribute, statements are made about *Moby_Dick*. Then in the second R element, pointing to the first R element using the *rpv* attribute, the statements are stated to be not trustable.

Query 38. Reification in RPV.

```
<R id="foo" r="http://example.org/Moby_Dick">
  <PV p="http://xmlns.com/foaf/0.1/maker" v="http://example.org/Oliver" />
</R>
<R rpv="#foo">
  <PV p="http://example.org/trustable" v="http://example.org/false" />
</R>
```

Though never actually used in any implementation, RPV animated others to pursue the goal of a XML-based format that emphasizes the triples instead of trying to somehow map RDF graphs to XML trees.

A.5.2 TriX

Following up on the requirements proposed by Beckett, Jeremy J. Carroll and Patrick Stickler in [66] designed the XML-based format TriX. The authors take a unique approach by first defining an absolutely minimal base format without any abbreviations and then using XSLT for syntactic extensions, together with stylesheets that convert to the base format. In addition, TriX goes beyond the original RDF model by supporting Named Graphs [65] and literals as subjects.

A TriX document contains one or more graphs, each optionally with a name. Each graph consists of one or more triples. The `triple` element is the core of TriX, containing three children. The position of each child determines whether the child is the subject, the property or the object of the triple. The element used identifies its type. The `uri` element is used for unabbreviated URIs, while the `id` element is used for identifying blank nodes. `plainLiteral` is used for String literals, while `typedLiteral` is used for any other type of literal in combination with a `datatype` attribute.

Query 39. Figure A.1 in TriX.

```
<TriX xmlns="http://www.w3.org/2004/03/trix/trix-1/">
  <graph>
    <uri>http://www.fakeroot.net/sw/SampleG</uri>
    <triple>
      <uri>http://www.fakeroot.net/sw/rdf-formats/</uri>
      <uri>http://www.w3.org/1999/02/22-rdf-syntax-ns#type</uri>
      <uri>http://xmlns.org/foaf/0.1/Document</uri>
    </triple>
    <triple>
      <uri>http://www.fakeroot.net/sw/rdf-formats/</uri>
      <uri>http://purl.org/dc/elements/1.1/title</uri>
      <plainLiteral>A Brief History of RDF Serialization Formats</plainLiteral>
    </triple>
    <triple>
      <uri>http://www.fakeroot.net/sw/rdf-formats/</uri>
      <uri>http://xmlns.org/foaf/0.1/maker</uri>
      <id>x</id>
    </triple>
    <triple>
      <id>x</id>
      <uri>http://xmlns.org/foaf/0.1/weblog</uri>
      <uri>http://slashdot.jp/~oliver/journal/</uri>
    </triple>
    <triple>
      <id>x</id>
      <uri>http://xmlns.org/foaf/0.1/name</uri>
      <plainLiteral>Oliver M. Bolzer</plainLiteral>
    </triple>
  </graph>
</TriX>
```

Example 39 is the sample RDF graph serialized using the basic TriX format. The name `http://www.fakeroot.net/sw/SampleG` is attached to it via the `uri` element directly under `graph`. Though being very verbose, the triples are clearly identifiable.

TriX allows syntactic extensions that make the syntax more human-friendly through the use of XSLT. One popular trick to increase readability of RDF serializations is to allow a XML QName-like abbreviation for URIs. By declaring an appropriate stylesheet processing instruction, TriX allows such syntactic sugar. Example 40 is one triple from the graph serialized in

TriX, using the `qname` element to abbreviate long URIs.

Query 40. Syntactic Extentions in TriX using XSLT.

```
<?xml-stylesheet type="text/xml"
  href="http://www.w3.org/2004/03/trix/all.xsl" ?>
<TriX xmlns="http://www.w3.org/2004/03/trix/trix-1/"
  xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
  xmlns:foaf="http://xmlns.org/foaf/0.1/" >
  <graph>
    <triple>
      <uri>http://www.fakeroot.net/sw/rdf-formats/</uri>
      <qname>rdf:type</qname>
      <qname>foaf:Document</qname>
    </triple>
  </graph>
</TriX>
```

Other syntactic sugars demonstrated by the authors of TriX include the use of `xml:base` as another method for URI-abbreviation, tags for specific typed literals and collections. The authors even go as far as suggesting RDF/XML as an TriX extention, based on the possibility of writing an RDF/XML parser in XSLT.

A.5.3 RXR

Discontent with TriX's decision to support features beyond the original RDF model, it's dependency on XSLT and the risk of ad-hoc extentions, Dave Beckett continued with the work he had began in [25] and formulated in [15] another proposal for a triple-centric XML-based format, RXR (Regular XML RDF).

Similar to TriX, the `triple` element, containing three children, is at the heart of RXR. But, instead of relying on the position, the children's roles are unambiguously identified by the elements `subject`, `predicate` and `object`. URIs are then given as value to the `uri` attribute, while literals are given as element content, with an optional `datatype` attribute. Blank nodes are specified with the `blank` attribute.

Query 41. Figure A.1 in RXR.

```
<graph xmlns="http://ilrt.org/discovery/2004/03/rxr/">
  <triple>
    <subject uri="http://www.fakeroot.net/sw/rdf-formats/" />
    <predicate uri="http://www.w3.org/1999/02/22-rdf-syntax-ns#type" />
    <object uri="http://xmlns.org/foaf/0.1/Document" />
  </triple>
  <triple>
    <subject uri="http://www.fakeroot.net/sw/rdf-formats/" />
    <predicate uri="http://purl.org/dc/elements/1.1/title" />
    <object>A Brief History of RDF Serialization Formats</object>
  </triple>
  <triple>
    <subject uri="http://www.fakeroot.net/sw/rdf-formats/" />
    <predicate uri="http://xmlns.org/foaf/0.1/maker" />
    <object blank="x" />
  </triple>
  <triple>
    <subject blank="x" />
    <predicate uri="http://xmlns.org/foaf/0.1/weblog" />
    <object uri="http://slashdot.jp/~oliver/journal/" />
  </triple>
```

```

</triple>
<triple>
  <subject blank="x" />
  <predicate uri="http://xmlns.org/foaf/0.1/name" />
  <object>Oliver M. Bolzer</object>
</triple>
</graph>

```

Example 41 shows again the RDF graph from Figure 1, this time serialized using RXR. Despite its verbosity, the triples are clearly recognizable.

RXR does not allow any abbreviations of URIs or other complexities such as XML literals. One notable exception are collections, supported by RXR through the `collection` element. Multiple statements with the same subject and predicate can be aggregated using this facility.

Query 42. Collections in RXR.

```

<graph xmlns="http://ilrt.org/discovery/2004/03/rxr/">
  <triple>
    <subject uri="http://example.org/box" />
    <predicate uri="http://example.org/contains" />
    <collection>
      <object>apple</object>
      <object>pear</object>
      <object>potato</object>
    </collection>
  </triple>
</graph>

```

Three triples are contained in Example 42. The same triples could also have been written separately, using three triple elements.

A.6 Features Overview

Feature\Format	RDF/XML	XMP	N3	N-Triples	Turtle	TriG	RPV	TriX basic	TriX ext.	RXR
XML	X	X					X	X	X	X
URI abbreviation	X	X	X		X	X	X		X	
Statement Aggregation	X	X	X		X	X	X			
Deep Nesting	X	X	X		X	X				
Blank Nodes	X	X	X	X	X	X	X	X	X	X
Collections	X	X	X		X	X			X	X
Typed Literals	X	X	X	X	X	X		X	X	X
Reification	X		X			X*	X	X*	X*	
Features beyond RDF			X			X		X	X	

*: Reification using Named Graphs

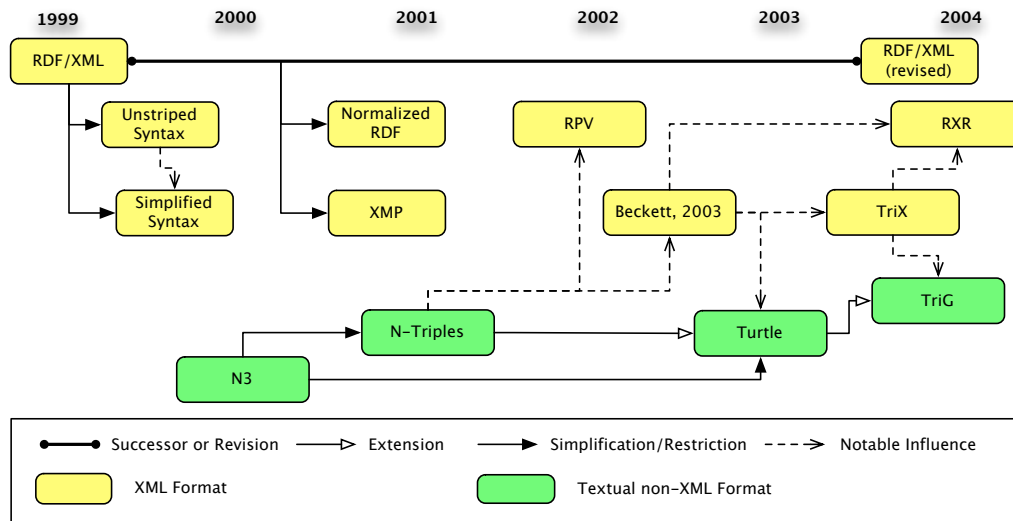
Normalized RDF, the Unstriped Syntax and Simplified Syntax are not listed, as they were only incomplete sketches and not concrete format proposals.

A.7 Genealogy

Figure A.2 is an attempt at portraying the genealogy of the formats described in this document. The newer formats were obviously influenced by most if not all preceding formats. Yet some

had a stronger influence, others less. The depicted affinities are based on citations by the proposals and statements made by the respective authors.

Figure A.2 Genealogy of RDF serialization formats
 (“Beckett, 2003” refers to [25]. Though not a concrete format proposal itself, it is included here as a major source of inspiration and guidance for the subsequent format proposals.)



A.8 Conclusions

The challenge of serializing RDF graphs and the dissatisfaction of the Semantic Web community with RDF/XML has brought forward numerous proposals for alternate serialization formats, most of which have been described here. After various early attempts to simplify RDF/XML failed to gain support, the idea of directly mapping RDF nodes and edges to XML elements appears to have been abandoned in favor of a more triple-centric view of RDF graphs.

N3 has seen wide adoption by the Semantic Web community as a triple-centric and human-friendly format. However actual implementations vary greatly in the supported features of N3. Current developments indicate a high chance that future implementations will standardize on Turtle as an adequate common denominator of the N3-based proposals.

Still, many in the community feel the need for a triple-centric XML-based format, in order to facilitate interchange between heterogeneous systems leveraging existing XML tools. TriX and RXR are the current contestants for such a format, but it is still too early to speculate on which will prevail.

Considering the disputes concerning support for new features like graph naming and literals as subjects, it is likely that the world will see yet more format proposals in the near future. Until some consensus is reached, RDF/XML remains the only formally standardized format all implementations must support.

Appendix B

Evaluation Tables

For space reasons, the detailed evaluation results are given for only a limited selection of languages. The full details can be obtained from the authors and will soon be available on the working group page: <http://rewerse.net/I4>.

Table B.1: Evaluation Table: Ease of Use

	1.1	1.2	1.3	1.4	1.5	1.6	1.7	1.8	1.9	1.10	1.11	1.12
XQL	•	-										
XPath 1.0 [78]	•	- 1	-	-	-	- ²	◦ ³	XML	single tree ⁴	regular path expressions, Unix-like paths	-	paths ⁵
XSLT 1.0 [77]	◦ ⁶	- 7	-	•	◦ ⁸	module-like inclusion mechanism, templates ⁹	◦	XML, [RDF ¹⁰]	single tree ⁴	stylesheet languages	functional	paths ¹¹
XPath 2.0 [30]	•	-	-	- ¹²	-	- ²	◦ ³	XML	single tree ¹³	regular path expressions, Unix-like paths	-	paths ⁵
XSLT 2.0 [150]	◦ ⁶	-	-	•	◦ ⁸	module-like inclusion mechanism, templates	•	XML	multiple trees ¹³	stylesheet languages	functional	paths ¹⁴

¹ There are a number of tools providing visual support for *authoring* XPath queries. ² As XPath is first and foremost intended to be used within a host language, it makes provisions for functions externally defined in the host language. ³ The evaluation context of an XPath expression can contain a library of arbitrary *externally defined* functions. ⁴ By means of the `id` function graph-like navigation is supported. ⁵ XPath predicates enable the specification of tree-like queries. ⁶ XSLT itself is specified using XML, leading to a rather verbose syntax but uses XPath for pattern specifications, thereby allowing programs to remain at least somewhat readable for humans. ⁷ There are a number of tools providing visual support for *authoring* XSLT programs. ⁸ As XSLT uses XPath for node access in various places, considerable effort is required to manipulate XSLT programs within XSLT. ⁹ Templates are limited in XSLT 1.0 w.r.t. to modularity, as any constructed XML data (i.e., that is not entirely extracted from the source document) has a special type that can not be queried further (cf. [169] and [150] where this problem has been addressed). ¹⁰ Using some form of extension functions as provided, e.g., by TreeHugger [245]. ¹¹ Uses XPath 1.0 for node access. ¹² A subset of the syntax for XQuery 1.0 specified in [177] could be used as an XML syntax for XPath 2.0. ¹³ By means of the `fn:id` and `fn:idref` functions graph-like navigation is supported. ¹⁴ Uses XPath 2.0 for node access.

Table B.1: Evaluation Table: Ease of Use

	1.1	1.2	1.3	1.4	1.5	1.6	1.7	1.8	1.9	1.10	1.11	1.12
XQuery 1.0 [41]	•	-	-	• ¹⁵	• ¹⁶	(optional) support for modules, functions	•	XML, [RDF, Topic Maps ¹⁷]	multiple trees ¹³	SQL, iterations in imperative languages	functional	paths
SquishQL [194]	•	-	-	-	-	-	-	RDF	triple	SQL	-	steps
rdfDB [128]												
RDQL [232]	•	-	-	-	-	-	-	RDF, RDFS ¹⁸	triple	SQL	-	steps
BRQL [220]	•	-	-	-	-	-	◦ ¹⁹	RDF	triple	SQL	-	steps
TriQL [40]	•	-	-	-	-	-	-	RDF	triple ²⁰	SQL	-	steps
“Syntactic Web” [226]	•	-	-	◦ ¹⁵	• ¹⁶	(optional) support for modules, functions	•	XML, RDF, Topic Maps	multiple trees ¹³	SQL, iterations in imperative languages	functional	paths
XsRQL [149]	•	-	-	-	-	- ²¹	- ²¹	RDF	graph ²²	XQuery	functional	paths
TreeHugger [245]	◦	-	-	•	◦	see XSLT 1.0	◦	RDF, XML	unclear, likely multiple trees	see XSLT 1.0	functional	paths
RDF Twig [257]	◦	-	-	•	◦	see XSLT 1.0	◦	RDF, XML	multiple trees	see XSLT 1.0	functional	paths
RDFT [88]	◦	-	-	•	◦	see XSLT 1.0	◦	RDF	unclear	see XSLT 1.0	functional	paths
Nexus [85]												
RDFPath [208]	•	-	-	-	-	-	-	RDF	single graph	XPath	-	paths ²³
RPath [185]	•	-	-	-	-	-	-	RDF	single acyclic graph	XPath	-	paths ²³

Table B.2: Evaluation Table: Functionality—Supported Query Types

	2.1.1	2.1.2	2.1.3	2.1.4	2.1.5	2.1.6	2.1.7	2.1.8	2.1.9	2.1.10	2.1.11	2.1.12	2.1.13
XPath 1.0 [78]	◦	set	-	•	-	•	-	•	•	-	-	◦ ²⁴	◦
XSLT 1.0 [77]	•	arbitrary XML or text	•	•	•	•	•	•	•	◦ ²⁵	•	•	◦

¹⁵ [177] specifies an XML syntax for XQuery 1.0 including path expressions. ¹⁶ However, XQueryX [177] is far from easy to process and there is a considerable conceptual gap for the programmer between the two syntactical forms, in particular where path expressions are concerned. ¹⁷ [224] proposes a canonical form for RDF and how to query that form using XQuery 1.0. ¹⁸ The implementation of RDQL in the Jena Toolkit [126] provides limited support for RDFS: the subsumption relations for classes and properties are treated (transparently) as transitive relations. ¹⁹ Additional functions on data values can be provided by an BRQL processor. However, no clear extension mechanism including the test whether some extension is supported has been defined so far. ²⁰ TriQL allows named graphs to be queried (cf. [65]). ²¹ User defined functions are planned to be added to the language. ²² The current draft is a bit unclear in this point. It only states that XsRQL builds upon an “RDF-oriented” data model. However the lack of closure and other recursive operations in XsRQL makes the issue of graph vs. tree data less critical. ²³ Predicates enable the specification of tree-like queries. ²⁴ Due to the lack of variables not all joins can be expressed. ²⁵ Some grouping queries can be expressed in XSLT 1.0, however more complex groupings tend to become extremely difficult to express.

Table B.2: Evaluation Table: Functionality—Supported Query Types

	2.1.1	2.1.2	2.1.3	2.1.4	2.1.5	2.1.6	2.1.7	2.1.8	2.1.9	2.1.10	2.1.11	2.1.12	2.1.13
XPath 2.0 [30]	◦	sequence	◦ ²⁶	•	–	•	–	•	•	–	–	•	◦
XSLT 2.0 [150]	•	arbitrary XML or text	•	•	•	•	•	•	•	•	•	•	•
XQuery 1.0 [41]	•	arbitrary XML or text	•	•	◦ ²⁷	•	•	•	•	•	•	•	•
SquishQL [194]	•	[table]	•	◦ ²⁸	◦ ²⁸	–	–	–	–	–	–	–	–
RDQL [232]	•	[table]	•	◦ ²⁸	◦ ²⁸	–	–	–	–	–	–	–	– ²⁹
BRQL [220]	•	[table, RDF triples]	•	◦ ²⁸	◦ ²⁸	•	◦	–	–	–	–	•	◦ ³⁰
TriQL [40]	•	[table]	•	◦ ²⁸	◦ ²⁸	–	–	–	–	–	–	–	–
“Syntactic Web” [226]	•	arbitrary XML or text	•	•	◦ ²⁷	•	•	•	•	•	•	•	•
XsRQL [149]	•	any text, [triples]	•	•	– ³¹	◦ ³²	– ³³	–	◦ ³⁴	–	◦ ³⁵	◦	
TreeHugger [245]	•	arbitrary XML or text	•	•	•	•	•	•	•	◦ ²⁵	•	•	◦
RDF Twig [257]	•	arbitrary XML or text	•	•	•	•	•	•	•	◦ ²⁵	•	•	◦
RDFT [88]	◦	?	◦ ³⁶	•	•	–	?	–	–	–	–	◦	–
RDFPath [208]	◦	[sets (of nodes or edges in the RDF graph)]	–	–	–	–	–	–	–	–	–	–	–
RPath [185]	◦	[sets (of nodes or edges in the RDF graph)]	– ³⁷	◦	–	–	–	–	–	–	–	–	–

²⁶ Due to the lack of construction and the fact, that XPath 2.0 always returns a single (flat) sequence, the selection of related information is limited to simple cases or use within a larger query. ²⁷ Early drafts of XQuery 1.0 contained a specialized function called `filter()` that allowed the reduction of the input structure to interesting elements while retaining their structural relations where possible. This function has been dropped, as similar functionality can be achieved by writing a recursive function as described in Appendix G.4 of [41]. ²⁸ Only substructures of previously known extent can be extracted or left out during reduction. ²⁹ The implementation of RDQL in the Jena Toolkit [126] provides limited support for inference: the RDFS subsumption relations for classes and properties are treated (transparently) as transitive relations. ³⁰ Only non-recursive inference is supported. ³¹ User defined and possibly recursive functions for structure traversal are planned to be added to the language. ³² Only test for not-equality is supported, but no general negation, e.g., for testing the non-existence of a property. ³³ Output format seems to be text-only in the current draft. ³⁴ An XPath-like `count()` function is provided, that counts the number of elements in a sequence, but no aggregation over the resulting values (such as in “return the book with the largest amount of authors”) is possible yet. ³⁵ Only sorting on the values of a sequence, no multi-key sorting. ³⁶ No support for variables in the current draft. ³⁷ [185] is unclear on the use of path expressions inside predicates. Only with paths expressions inside predicates related information can be queried.

Table B.3: Evaluation Table: Adequacy—XML

	2.2.1	2.2.2	2.2.3	2.2.4	2.2.5	2.2.6	2.2.7	2.2.8	2.2.9	2.2.10	2.2.11	2.2.12	2.2.13	2.2.14	2.2.15
Table B.3: Evaluation Table: Adequacy—XML															
	2.2.1	2.2.2	2.2.3	2.2.4	2.2.5	2.2.6	2.2.7	2.2.8	2.2.9	2.2.10	2.2.11	2.2.12	2.2.13	2.2.14	2.2.15
XPath 1.0 [78]	•	•	•	explicit	•	◦ ³⁸	•	•	– ₃₉	– ₄₀	–	–	◦ ⁴¹	–	–
XSLT 1.0 [77]	•	•	•	explicit	•	◦ ³⁸	•	•	– ₃₉	– ₄₀	•	–	•	– ₄₂	– ₄₃
XPath 2.0 [30]	•	•	•	explicit	•	◦ ³⁸	•	•	– ₃₉	– ₄₀	◦ ⁴⁴	•	◦ ⁴¹	–	–
XSLT 2.0 [150]	•	•	•	explicit	•	◦ ³⁸	•	•	– ₃₉	– ₄₀	•	•	•	–	–
XQuery 1.0 [41]	•	•	•	explicit	•	◦ ³⁸	•	•	– ₃₉	– ₄₀	•	•	•	–	–
“Syntactic Web” [226]		see XQuery 1.0													
TreeHugger [245]		see XSLT 1.0													
RDF Twig [257]		see XSLT 1.0													

Table B.4: Evaluation Table: Adequacy—RDF

	2.2.5	2.2.10	2.3.1	2.3.2	2.3.3	2.3.4	2.3.5	2.3.6	2.3.7	
XSLT 1.0 [77]		see TreeHugger and RDF Twig								
XQuery 1.0 [41]		see “Syntactic Web”								
SquishQL [194]	–	–	•	•	–	–	–	–	–	
RDQL [232]	–	–	•	•	–	–	–	–	–	
BRQL [220]	–	–	•	•	–	–	◦ ⁴⁵	•	–	
TriQL [40]	–	–	•	•	–	–	◦ ⁴⁶	–	–	
“Syntactic Web” [226]	•	– ₄₀	•	•	–	–	–	–	any	
XsRQL [149]	–	–	•	•	–	–	◦ ⁴⁷	◦	any text, [triples]	

³⁸ Predicates allow the specification of several children of a single node in the XML structure without giving their order amongst each other. However, enforcing that these children are different requires the use of variables and complicates the query noticeably. ³⁹ By a combination of the XPath constructs `following-sibling`, `preceding-sibling` and `count()` it is possible to simulate total queries. However, authoring such queries is very cumbersome. ⁴⁰ Using `if-then-else` or `union` constructs allows the specification of alternative queries of which one includes the optional part and the other omits it, therefore ensuring (barring duplicate removal) that the optional part is returned where possible. ⁴¹ XPath can query the namespace of an element, however namespaces prefixes have to be set-up outside of XPath. ⁴² Some XSLT processors provide extension functions for evaluating an XPath expression contained in a string (e.g., from an attribute value). But none of these extensions considers the extended functionality of XPointer. ⁴³ Some XSLT processors provide extension functions for outputting multiple documents, thus enabling the treatment of compound result documents. ⁴⁴ Construction in XPath 2.0 is limited to constructing flat sequences. ⁴⁵ Support for “quads” is provided, hence sometimes querying the reified form of a statement is unnecessary. ⁴⁶ TriQL supports the use of “named graphs” as described in [65]. Named graphs can be seen as a form of reification. Only this form of reification is supported in TriQL, statements reified by the usual RDF convention are not treated specifically. ⁴⁷ XsRQL provides some special functions for treating quads.

Table B.4: Evaluation Table: Adequacy—RDF

	2.2.5	2.2.10	2.3.1	2.3.2	2.3.3	2.3.4	2.3.5	2.3.6	2.3.7
TreeHugger [245]	•	⁴⁰ –	•	•	⁴⁸ ◦	–	–	–	any
RDF Twig [257]	•	⁴⁰ –	•	•	–	–	–	–	any
RDFT [88]	–	–	•	•	–	–	–	–	?
RDFPath [208]	–	–	•	–	–	–	–	–	[sets (of nodes or edges in the RDF graph)]
RPath [185]	⁴⁹ ◦	–	•	–	•	–	–	–	[sets (of nodes or edges in the RDF graph)]

Table B.5: Evaluation Table: Adequacy—Topic Maps

	2.2.10	2.3.2	2.4.1	2.4.2	2.4.3	2.4.4	2.4.5	2.4.6	2.4.7
XQuery 1.0 [41]		see “Syntactic Web”							
“Syntactic Web” [226]	⁴⁰ –	•	•	•	•	◦	–	–	any

Table B.6: Evaluation Table: Adequacy—Web

	2.5.1	2.5.2	2.5.3	2.5.4	2.5.5	2.5.6	2.5.7	
XPath 1.0 [78]	–	–	–	–	◦	special	–	
XSLT 1.0 [77]	⁵⁰ –	⁵¹ –	–	–	⁵² ◦	text, XML, [RDF, Topic Maps]	•	
XPath 2.0 [30]	•	–	–	⁵³ ◦	◦	special	–	
XSLT 2.0 [150]	•	•	–	⁵³ ◦	⁵⁴ ◦	text, XML, [RDF, Topic Maps]	•	
XQuery 1.0 [41]	•	•	–	⁵³ ◦	⁵⁵ ◦	text, XML, [RDF, Topic Maps]	•	
SquishQL [194]	•	–	–	–	–	[variable bindings]	–	
RDQL [232]	•	–	–	–	–	[variable bindings]	–	
BRQL [220]	•	–	•	–	–	[variable bindings, RDF triples]	–	
TriQL [40]	•	–	–	–	–	[variable bindings]	–	
“Syntactic Web” [226]		see XQuery 1.0						
XsRQL [149]	–	–	–	–	⁵⁶ ◦	any text, [triples]	◦	
TreeHugger [245]		see XSLT 1.0						
RDF Twig [257]		see XSLT 1.0						
RDFT [88]	–	–	–	–	?	?	–	
RDFPath [208]	–	–	–	–	–	[sets (of nodes or edges in the RDF graph)]	–	

⁴⁸ Based upon the XML document order, elements of a container can be conveniently queried. However, it is unclear from the specification, whether this covers all cases, e.g., when containers are stated explicitly using `rdf:_1`, etc.

⁴⁹ A function for accessing the *i*-th item in a sequence container is provided. ⁵⁰ Some XSLT 1.0 processors provide extension functions for accessing multiple source documents. ⁵¹ Some XSLT 1.0 processors provide extension functions for outputting to multiple targets. ⁵² If the output is XML, queries can be composed, otherwise not. ⁵³ Multiple documents are queried through the use of the `fn:doc()` function. It is possible to identify which queries are to be evaluated on which source, but the query language allows a close intertwining of queries to different documents. ⁵⁴ If the output is XML, queries can be composed, otherwise not. ⁵⁵ If the output is XML, queries can be composed, otherwise not. ⁵⁶ If the output are triples or quads, it can be further processed by other XsRQL queries.

Table B.6: Evaluation Table: Adequacy—Web

	2.5.1	2.5.2	2.5.3	2.5.4	2.5.5	2.5.6	2.5.7
RPath [185]	-	-	-	-	-	[sets (of nodes or edges in the RDF graph)]	-

Table B.7: Evaluation Table: Adequacy—Approximate Answering and Full-text Processing

	2.6.1	2.6.2	2.6.3	2.6.4	2.6.5
XPath 1.0 [78]	-	-	◦ ⁵⁷	◦ ⁵⁸	-
XSLT 1.0 [77]	-	-	◦ ⁵⁷	◦ ⁵⁸	-
XPath 2.0 [30] ⁵⁹	•	•	•	•	•
XSLT 2.0 [150]	•	◦ ⁶⁰	•	•	•
XQuery 1.0 [41] ⁵⁹	•	•	•	•	•
SquishQL [194]	-	-	-	-	-
RDQL [232]	-	-	-	-	-
BRQL [220]	-	-	-	◦ ⁶¹	-
TriQL [40]	-	-	-	-	-
“Syntactic Web” [226]	see XQuery 1.0				
XsRQL [149]	-	-	-	-	-
TreeHugger [245]	see XSLT 1.0				
RDF Twig [257]	see XSLT 1.0				
RDFT [88]	-	-	-	-	-
RDFPath [208]	-	-	-	-	-
RPath [185]	-	-	-	-	-

Table B.8: Evaluation Table: Functionality—Evolution and Reactivity

	2.7.1	2.7.2	2.7.3
XPath 1.0 [78]	-	-	-
XSLT 1.0 [77]	◦ ⁶²	-	◦ ⁶²
XPath 2.0 [30]	-	-	-
XSLT 2.0 [150]	◦ ⁶²	-	◦ ⁶²
XQuery 1.0 [41]	◦ ⁶³	-	◦ ⁶⁴
SquishQL [194]	-	-	-
RDQL [232]	◦ ⁶⁵	-	-
BRQL [220]	-	-	-
TriQL [40]	-	-	-
“Syntactic Web” [226]	see XQuery 1.0		
XsRQL [149]	-	-	-
TreeHugger [245]	see XSLT 1.0		

⁵⁷ Partial queries can be seen as some form of approximate structure matching. ⁵⁸ It is possible to test, whether a word occurs inside an element. ⁵⁹ This applies only, if [12] is in use. Otherwise, the evaluation becomes similar to the one for XPath 1.0. ⁶⁰ Although the XPath 2.0 full-text extension [12] provide scoring of results there is no support for this on the level of XSLT, e.g., for constructing several ranked variants. ⁶¹ Through the use of functions on data values some kind of full-text processing could be provided. No such provision has been made so far. ⁶² Standard XSLT does not support updates or reactivity, however [47] proposes a reactive extension for XSLT (and Lorel). ⁶³ Several extensions for XQuery have been proposed that provide update capabilities, e.g., XUpdate [168] and [248]. ⁶⁴ Several extensions for XQuery have been proposed that provide reactive behavior, e.g., [45]. ⁶⁵ RIDIQL [259] extends RDQL by simple update commands.

Table B.8: Evaluation Table: Functionality—Evolution and Reactivity

	2.7.1	2.7.2	2.7.3
RDF Twig [257]		see XSLT 1.0	
RDFT [88]	-	-	-
RDFPath [208]	-	-	-
RPath [185]	-	-	-

Table B.9: Evaluation Table: Semantics

	3.1	3.2	3.3	3.4	3.5
XPath 1.0 [78]	◦ ⁶⁶	•	•	◦ ⁶⁷	-
XSLT 1.0 [77]	◦ ⁶⁸	•	•	◦ ⁶⁹	• ⁷⁰
XPath 2.0 [30]	•	•	•	•	-
XSLT 2.0 [150]	-	N/A	•	-	-
XQuery 1.0 [41]	•	•	•	•	-
SquishQL [194]	-	N/A	•	-	-
RDQL [232]	-	N/A	•	-	-
BRQL [220]	-	N/A	•	-	-
TriQL [40]	-	N/A	•	-	-
"Syntactic Web" [226]		see XQuery 1.0			
XsRQL [149]					
TreeHugger [245]	-	N/A	•	-	-
RDF Twig [257]	-	N/A	•	-	-
RDFT [88]	-	N/A	•	-	-
RDFPath [208]	-	N/A	•	-	-
RPath [185]	-	N/A	•	-	-

Table B.10: Evaluation Table: Formal Properties and Implementation

	4.1	4.2	4.3	4.4	4.5	4.6	4.7	4.8	4.9	4.10	4.11	4.12
XPath 1.0 [78]	•	-	• ⁷¹	• ⁷²	◦ ⁷³	numerous	production use	full	•	•	• ⁷⁴	• ⁷⁴
XSLT 1.0 [77]	•	• ⁷⁵	◦ ⁷⁶	-	-	numerous	production use	full	-	•	◦	• ⁷⁷
XPath 2.0 [30]	•	-	◦ ⁷⁸	• ⁷⁸	◦ ⁷⁸	numerous ⁷⁹	wide-spread	full ⁷⁹	•	• ⁷⁹	• ⁷⁹	
XSLT 2.0 [150]	•	• ⁷⁵	◦ ⁷⁶	-	-	few	prototype	few	-	•	-	• ⁸⁰
XQuery 1.0 [41]	•	• ⁷⁵	◦ ⁷⁸	• ⁷⁸	-	numerous	prototype to production use	full	•	•	• ⁸¹	• ⁸²
SquishQL [194]	•	no	-	-	-	8 ⁸³	production use	wide-spread	•	•	-	-

⁶⁶ Although no formal semantics for XPath 1.0 has been published by the W3C, there are several proposals for a formal semantics [253, 254, 207] that cover a large subset of XPath 1.0. ⁶⁷ Both [254, 207] propose formal semantics that can also be seen as operational semantics for a (naive) functional implementation. ⁶⁸ Although no formal semantics for XSLT 1.0 has been published by the W3C, there are several proposals for formal semantics of language subsets [253, 153], the most complete being [153]. ⁶⁹ Where a formal semantics for XSLT 1.0 has been proposed, it can be easily used as basis for an implementation, as demonstrated in [153]. ⁷⁰ An abstract (or "virtual") machine for XSLT has been described in [198]. ⁷¹ See [121, 28]. ⁷² See [119, 122, 120, 121]. ⁷³ See [28]. ⁷⁴ See, e.g., [205, 206]. ⁷⁵ See [151]. ⁷⁶ E.g., "core XSLT" in [153]. ⁷⁷ E.g., Saxon <http://saxon.sourceforge.net/>. ⁷⁸ The properties of XPath 1.0 can be inherited here. ⁷⁹ Via XQuery 1.0 implementations. ⁸⁰ E.g., Saxon <http://saxon.sourceforge.net/>. ⁸¹ See [99, 106, 157]. ⁸² See [99, 106, 157]. ⁸³ However, 6 of these actually implement RDQL.

Table B.10: Evaluation Table: Formal Properties and Implementation

	4.1	4.2	4.3	4.4	4.5	4.6	4.7	4.8	4.9	4.10	4.11	4.12
RDQL [232]	•	no	-	-	-	6	production use	wide-spread	•	•	-	-
BRQL [220]	•	no	-	-	-	0	N/A	N/A	-	-	-	-
TriQL [40]	•	no	-	-	-	0	N/A	N/A	-	-	-	-
“Syntactic Web” [226]	•	•	-	⁸⁴	-	⁸⁵	N/A	N/A	N/A	N/A	N/A	N/A
XsRQL [149]	•	unclear, but unlikely	-	-	-	0	N/A	N/A	N/A	N/A	N/A	N/A
TreeHugger [245]	•	•	-	-	-	1	prototype	-	-	•	-	-
RDF Twig [257]	•	•	-	-	-	1	prototype	-	-	•	-	-
RDFT [88]	•	unclear, but unlikely	-	-	-	0	N/A	N/A	N/A	N/A	N/A	N/A
RDFPath [208]	•	-	-	-	-	0	N/A	N/A	N/A	N/A	N/A	N/A
RPath [185]	•	-	-	-	-	1	prototype	-	-	•	-	-

Table B.11: Evaluation Table: Reasoning

	5.1	5.2	5.3	5.4	5.5	5.6	5.7	5.8	5.9	5.10	
XPath 1.0 [78]	-	•	•	•	⁸⁶	-	•	-	-	-	
XSLT 1.0 [77]	⁸⁷	•	•	•	⁸⁶	•	•	-	-	-	
XPath 2.0 [30]	-	•	•	•	•	-	•	-	-	-	
XSLT 2.0 [150]	•	•	•	•	•	•	•	-	-	-	
XQuery 1.0 [41]	•	•	•	•	•	•	•	-	-	-	
SquishQL [194]	-	•	•	⁸⁸	-	-	-	-	N/A	N/A	
RDQL [232]	-	•	•	⁸⁸	-	-	-	⁸⁹	N/A	N/A	
BRQL [220]	-	•	•	•	-	-	-	-	N/A	N/A	
TriQL [40]	-	•	•	⁸⁸	-	-	-	-	N/A	N/A	
“Syntactic Web” [226]	-	see XQuery 1.0									
XsRQL [149]	-	•	•	-	⁹⁰	-	-	-	-	-	
TreeHugger [245]	-	see XSLT 1.0									
RDF Twig [257]	-	see XSLT 1.0									
RDFT [88]	-	-	-	-	-	-	-	-	-	-	
RDFPath [208]	-	•	•	-	-	-	-	-	-	-	
RPath [185]	-	-	-	-	-	-	•	-	-	-	

⁸⁴ The polynomial core for XPath (and therefore for XQuery) does naturally not contain function definitions, let alone recursive ones as used by this proposal. ⁸⁵ But could be implemented on most of the numerous XQuery implementations. ⁸⁶ Through the use of the XPath count() function it is possible to emulate all quantification. ⁸⁷ XSLT 1.0 templates can provide views on the data, however limited by the inability to query any structure constructed (and not extracted from the input). ⁸⁸ No test whether a tuple does not exist. ⁸⁹ The Jena Toolkit [95] implementation of RDQL provides the transitive closure over the RDFS sub-class and sub-property relations. The RDF extension RIDIQL [259] provides means to use different reasoners for extending a bare RDF graph based on the semantics for RDFS or OWL. ⁹⁰ XsRQL provides a XPath-style count() function that can be used to implement all-quantification.

Table B.12: Evaluation Table: Ontology Awareness

	6.1.1	6.2.1	6.2.2	6.2.3	6.3.1	6.3.2	6.3.3	6.3.4	6.3.5
XPath 1.0 [78]	-	-	N/A	N/A	-	N/A	N/A	N/A	N/A
XSLT 1.0 [77]			see TreeHuger and RDF Twig						
XPath 2.0 [30]	-	-	N/A	N/A	-	N/A	N/A	N/A	N/A
XSLT 2.0 [150]			see TreeHuger and RDF Twig						
XQuery 1.0 [41]			see “Syntactic Web”						
SquishQL [194]	-	-	N/A	N/A	-	N/A	N/A	N/A	N/A
RDQL [232]	◦ ⁸⁹	◦ ⁸⁹	◦ ⁸⁹	-	◦ ⁸⁹	◦ ⁸⁹	◦ ⁸⁹	◦ ⁸⁹	-
BRQL [220]	-	-	N/A	N/A	-	N/A	N/A	N/A	N/A
TriQL [40]	-	-	N/A	N/A	-	N/A	N/A	N/A	N/A
“Syntactic Web” [226]	◦	◦	explicit	-	-	N/A	N/A	N/A	N/A
XsRQL [149]	-	-	N/A	N/A	-	N/A	N/A	N/A	N/A
TreeHuger [245]	◦	◦	•	◦ ⁹¹	◦	•	unclear	unclear	
RDF Twig [257]	-	-	N/A	N/A	-	N/A	N/A	N/A	N/A
RDFT [88]	◦	◦	transparent	-	-	N/A	N/A	N/A	N/A
RDFPath [208]	-	-	N/A	N/A	-	N/A	N/A	N/A	N/A
RPath [185]	-	-	N/A	N/A	-	N/A	N/A	N/A	N/A

Table B.13: Evaluation Table: Ontology Awareness—Typing

	6.4.1	6.4.2	6.4.3	6.4.4	6.4.5	6.4.6
XPath 1.0 [78]	untyped	N/A	N/A	N/A	N/A	-
XSLT 1.0 [77]	untyped	N/A	N/A	N/A	N/A	-
XPath 2.0 [30]	•	dynamic	both	-	•	•
XSLT 2.0 [150]	•	dynamic	both	-	•	•
XQuery 1.0 [41]	•	static	both	◦	•	•
SquishQL [194]	untyped	N/A	N/A	N/A	N/A	-
RDQL [232]	untyped	N/A	N/A	N/A	N/A	◦ ⁹²
BRQL [220]	◦	dynamic	explicit	-	-	•
TriQL [40]	◦	dynamic	explicit	-	-	•
“Syntactic Web” [226]	•	static	both	◦	•	•
XsRQL [149]	◦	dynamic	unclear	-	-	•
TreeHuger [245]	untyped	N/A	N/A	N/A	N/A	-
RDF Twig [257]	untyped	N/A	N/A	N/A	N/A	-
RDFT [88]	untyped	N/A	N/A	N/A	N/A	-
RDFPath [208]	untyped	N/A	N/A	N/A	N/A	-
RPath [185]	untyped	N/A	N/A	N/A	N/A	-

⁹¹ It is possible to query the extent of a class, however no static typing is provided. ⁹² Newer implementations support the use of XML Schema Datatypes in RDF.

Bibliography

- [1] iTQL Commands (online, retrieved at 2004-07-17). Available from: <http://www.kowari.org/271.htm>.
- [2] RDF Gateway (online, retrieved at 2004-07-17). Available from: <http://www.intellidimension.com/default.rsp?topic=/pages/rdfgateway/reference/db/default.rsp>.
- [3] RDFQL Database Command Reference (online, retrieved at 2004-07-17). Available from: <http://www.intellidimension.com/default.rsp?topic=/pages/rdfgateway/reference/db/default.rsp>.
- [4] S. Abiteboul. Querying Semi-Structured Data. In *Proc. International Conference on Database Theory*, 1997.
- [5] S. Abiteboul, D. Quass, J. McHugh, J. Widom, and J. Wiener. The Lorel Query Language for Semistructured Data. *International Journal on Digital Libraries*1(1):68-88, April 1997., 1(1):68-88, 1997. Available from: <http://www-db.stanford.edu/lore/pubs/Lore196.pdf>.
- [6] S. Abiteboul and V. Vianu. Queries and Computation on the Web. In *Proc. International Conference on Database Theory*, pages 262-275, 1997. Available from: <http://dbpubs.stanford.edu:8090/pub/1996-20>.
- [7] *Adobe Extensible Metadata Platform (XMP)*, 2001, Adobe Systems Inc. Available from: <http://www.adobe.com/products/xmp/>.
- [8] S. Alexaki, N. Athanasis, V. Christophides, G. Karvounarakis, A. Maganaraki, D. Plexousakis, and K. Toll. The ICS-FORTH RDFSuite: High-level Scalable Tools for the Semantic Web. *ERCIM News*, (51 (Special Issue on the Semantic Web)), 2002. Available from: http://www.ercim.org/publication/Ercim_News/enw51/alexaki.html.
- [9] S. Alexaki, V. Christophides, G. Karvounarakis, and D. Plexousakis. The ICS-FORTH RDFSuite: Managing Voluminous RDF Description Bases. In *Proc. International Workshop on the Semantic Web*, 2001. Available from: <http://athena.ics.forth.gr:9090/RDF/publications/semweb2001.pdf>.
- [10] J. Alferes, W. May, and P. Patranjan. *State of the Art on Evolution and Reactivity*, 2004, Deliverable.

- [11] G. Alliance. RDF::Core::Query—Implementation of Query Language (online). 2004 (retrieved at 2004-08-19). Available from: <http://www.gingerall.com/charlie/ga/html/rdf-core/RDF/Core/Query.pm.html>.
- [12] S. Amer-Yahia, C. Botev, S. Buxton, P. Case, J. Doerre, D. McBeath, M. Rys, and J. Shanmugasundaram, editors. *XQuery and XPath Full-Text*, 2004, W3C, Working Draft. Available from: <http://www.w3.org/TR/xquery-full-text-requirements/>.
- [13] C. Anutariya, V. Wuwongse, and V. Wattanapailin. An Equivalent-Transformation-Based XML Rule Language. In *Proc. International Workshop on Rule Markup Languages for Business Rules in the Semantic Web*, 2002. Available from: <http://kr.cs.ait.ac.th/old/publications/XETruleml.pdf>.
- [14] D. Backett. New Syntaxes for RDF (online). November 2003 (retrieved at 2004-07-09). Available from: <http://www.ilrt.bris.ac.uk/discovery/2003/11/new-syntaxes-rdf/>.
- [15] D. Backett. Modernising Semantic Web Markup. In *Proc. XML Europe*, April 2004. Available from: http://www.ideaalliance.org/papers/dx_xmle04/papers/03-08-03/03-08-03.html.
- [16] R. Barta. AsTma= Language Definition (online, retrieved at 2007-07-18). Available from: <http://astma.it.bond.edu.au/astma=-spec-xtm.dbk>.
- [17] R. Barta. AsTma? Tutorial. Technical report, Bond University, 2003. Available from: <http://astma.it.bond.edu.au/astma%3F-tutorial.dbk>.
- [18] R. Barta. Path Language for Topic Maps: Full speed ahead? (online). 2004 (retrieved at 2004-08-19). Available from: <http://topicmaps.it.bond.edu.au/docs/30>.
- [19] R. Barta and L. Garshol, editors. *Topic Map Query Language, Use Cases*, December 2003, ISO/IEC, Technical document. Available from: <http://www.y12.doe.gov/sgml/sc34/document/0449.htm>.
- [20] R. Barta and J. Gylta. *XTM::Path*, 2002. Available from: <http://cpan.uwinnipeg.ca/htdocs/XTM/XTM/Path.html>.
- [21] C. Baru, A. Gupta, B. Ludäscher, R. Marciano, Y. Papakonstantinou, and P. Velikhov. XML-Based Information Mediation with MIX. In *Proc. ACM SIGMOD International Conference on Management of Data*, 1999. Available from: <http://www.db.ucsd.edu/publications/vamp.pdf>.
- [22] N. Bassiliades and I. Vlahavas. Intelligent Querying of Web Documents Using a Deductive XML Repository. In *Proc. Hellenic Conference on Artificial Intelligence*, April 2002. Available from: <http://lpis.csd.auth.gr/publications/setn02-bassiliades.pdf>.
- [23] N. Bassiliades and I. Vlahavas. Capturing RDF Descriptive Semantics in an Object Oriented Knowledge Base System. In *Proc. International World Wide Web Conference*, May 2003. Available from: <http://www2003.org/cdrom/papers/poster/p277/p277-nbassili.html>.

- [24] S. Bechhofer, F. van Harmelen, J. Hendler, I. Horrocks, D. McGuinness, P. Patel-Schneider, and L. Stein. *OWL Web Ontology Language—Reference*, 2004, W3C, Recommendation. Available from: <http://www.w3.org/TR/owl-ref/>.
- [25] D. Beckett. A retrospective on the development of the RDF/XML Revised Syntax (online). June 2003 (retrieved at 2004-07-09). Available from: http://www.ilrt.bris.ac.uk/publications/researchreport/rr1017/report_html.
- [26] D. Beckett. *Turtle - Terse RDF Triple Language*, February 2004.
- [27] D. Beckett and B. McBride, editors. *RDF/XML Syntax Specification (Revised)*, 2004, W3C, Recommendation. Available from: <http://www.w3.org/TR/rdf-syntax-grammar/>.
- [28] M. Benedikt, W. Fan, and G. Kuper. Structural Properties of XPath Fragments. In *Proc. International Conference on Database Theory*, 2003. Available from: <http://homepages.inf.ed.ac.uk/wenfei/papers/icdt03.pdf>.
- [29] V. Benzaken, G. Castagna, and A. Frisch. CDuce: An XML-Centric General-Purpose Language. In *Proc. International Conference on Functional Programming*, 2003.
- [30] A. Berglund, S. Boag, D. Chamberlin, M. Fernandez, M. Kay, J. Robie, and J. Simeon, editors. *XML Path Language (XPath) 2.0*, 2003, W3C, Working Draft. Available from: <http://www.w3.org/TR/xpath20/>.
- [31] A. Berlea and H. Seidl. fxt—A Transformation Language for XML Documents. *Journal of Computing and Information Technology, Special Issue on Domain-Specific Languages*, 2001.
- [32] T. Berners-Lee. Notation 3, an RDF language for the Semantic Web (online, retrieved at 2004-07-09). Available from: <http://www.w3.org/DesignIssues/Notation3>.
- [33] T. Berners-Lee. Primer: Getting into RDF and Semantic Web using N3 (online, retrieved at 2004-07-09). Available from: <http://www.w3.org/2000/10/swap/Primer.html>.
- [34] T. Berners-Lee. Semantic Web Road Map (online, retrieved at 2004-08-09). Available from: <http://www.w3.org/DesignIssues/Semantic.html>.
- [35] T. Berners-Lee. A strawman Unstriped syntax for RDF in XML (online). May 1999 (retrieved at 2004-07-09). Available from: <http://www.w3.org/DesignIssues/Syntax.html>.
- [36] T. Berners-Lee. N3QL—RDF Data Query Language (online). 2004 (retrieved at 2004-07-17). Available from: <http://www.w3.org/DesignIssues/N3QL.html>.
- [37] T. Berners-Lee, J. Hendler, and O. Lassila. The Semantic Web—A new form of Web content that is meaningful to computers will unleash a revolution of new possibilities. *Scientific American*, 2001.
- [38] P. Biron and A. Malhotra, editors. *XML Schema Part 2: Datatypes*, 2001, W3C, Recommendation. Available from: <http://www.w3.org/TR/xmlschema-2/>.
- [39] C. Bizer. The TriG Syntax (online). April 2004 (retrieved at 2004-07-09). Available from: <http://www.wiwiss.fu-berlin.de/suhl/bizer/TriG/>.

- [40] C. Bizer. TriQL—A Query Language for Named Graphs (online). 2004 (retrieved at 2004-07-20). Available from: <http://www.wiwiss.fu-berlin.de/suhl/bizer/TriQL/Spec/>.
- [41] S. Boag, D. Chamberlin, M. Fernandez, D. Florescu, J. Robie, and J. Simeon, editors. *XQuery 1.0: An XML Query Language*, 2003, W3C, Working Draft. Available from: <http://www.w3.org/TR/xquery/>.
- [42] D. Bogachev. TPath - Introduction (online). 2003 (retrieved at 2004-04-27). Available from: <http://homepage.mac.com/dmitryv/TopicMaps/TPath/TPathIntroduction.html>.
- [43] D. Bogachev. TPath - Revisited (online). 2004 (retrieved at 2004-04-27). Available from: <http://homepage.mac.com/dmitryv/TopicMaps/TPath/TPathRevisited.html>.
- [44] H. Boley, B. Grosz, M. Sintek, S. Tabet, and G. Wagner. RuleML Design (online). 2002 (retrieved at 2004-08-09). Available from: <http://www.ruleml.org/indesign.html>.
- [45] A. Bonifati, D. Braga, A. Campi, and S. Ceri. Active XQuery. In *Proc. International Conference on Data Engineering*, 2002. Available from: <http://nike.psu.edu/classes/ist597/2003-fall/papers/ICDE02.pdf>.
- [46] A. Bonifati and S. Ceri. Comparative Analysis of Five XML Query Languages. *ACM SIGMOD Record*, 2000. Available from: <http://portal.acm.org/citation.cfm?id=344822>.
- [47] A. Bonifati, S. Ceri, and S. Paraboschi. Active Rules for XML: A New Paradigm for E-Services. *VLDB Journal*, 2001. Available from: <http://portal.acm.org/citation.cfm?id=767136&dl=ACM&coll=portal&CFID=11111111&CFTOKEN=2222222>.
- [48] A. Bonifati and D. Lee. Technical Survey of XML Schema and Query Languages. Technical report, University of California, Los Angeles, Computer Science Dept., 2001.
- [49] T. Bray. RPV: Triples Made Plain (online). 2002 (retrieved at 2004-07-09). Available from: <http://www.textuality.com/xml/RPV.html>.
- [50] T. Bray, D. Hollander, and A. Layman, editors. *Namespaces in XML*, 1999, W3C, Recommendation. Available from: <http://www.w3.org/TR/REC-xml-names/>.
- [51] T. Bray, D. Hollander, A. Layman, and R. Tobin, editors. *Namespaces in XML 1.1*, 2004, W3C, Recommendation. Available from: <http://www.w3.org/TR/xml-names11/>.
- [52] T. Bray, J. Paoli, C. M. Sperberg-McQueen, E. Maler, and F. Yergeau, editors. *Extensible Markup Language (XML) 1.0 (Third Edition)*, 2004, W3C, Recommendation. Available from: <http://www.w3.org/TR/REC-xml/>.
- [53] D. Brickley. RDF: Understanding the Striped RDF/XML Syntax (online). October 2001 (retrieved at 2004-07-09). Available from: <http://www.w3.org/2001/10/stripes/>.
- [54] D. Brickley, R. Guha, and B. McBride, editors. *RDF Vocabulary Description Language 1.0: RDF Schema*, 2004, W3C, Recommendation. Available from: <http://www.w3.org/TR/rdf-schema/>.

- [55] D. Brickley and L. Miller. *FOAF Vocabulary Specification*, May 2004. Available from: <http://xmlns.com/foaf/0.1/>.
- [56] J. Broekstra, C. Fluit, and F. van Harmelen. The State of the Art on Representation and Query Languages for Semistructured Data. On-To-Knowledge EU-IST-1999-10132 Deliverable 8, Administrator Nederland b.v., 2000. Available from: <http://www.ontoknowledge.org/countd/countdown.cgi?de18.pdf>.
- [57] J. Broekstra and A. Kampman. Query Language Definition. On-To-Knowledge EU-IST-1999-10132 Deliverable 9, Administrator Nederland b.v., 2001. Available from: <http://www.ontoknowledge.org/countd/countdown.cgi?de19.pdf>.
- [58] J. Broekstra and A. Kampman. SeRQL: A Second Generation RDF Query Language. In *Proc. SWAD-Europe Workshop on Semantic Web Storage and Retrieval*, 2003. Available from: <http://www.w3.org/2001/sw/Europe/events/20031113-storage/positions/aduna.pdf>.
- [59] J. Broekstra, A. Kampman, and F. Harmelen. Sesame: A Generic Architecture for Storing and Querying RDF and RDF Schema. In *Proc. International Semantic Web Conference*, 2002. Available from: <http://www.cs.vu.nl/~frankh/postscript/ISWC02.pdf>.
- [60] F. Bry, T. Furche, L. Badea, C. Koch, S. Schaffert, and S. Berger. Identification of Design Principles for a (Semantic) Web Query Language. Deliverable I4-D1, REVERSE, 2004. Available from: <http://reverse.net/I4/>.
- [61] P. Buneman, M. Fernandez, and D. Suciu. UnQL: A Query Language and Algebra for Semistructured Data Based on Structural Recursion. *VLDB Journal*, 9(1):76-110, 2000. Available from: <http://portal.acm.org/citation.cfm?id=765224&jmp=cit&coll=GUIDE&d1=ACM>.
- [62] S. Buxton and M. Rys, editors. *XQuery and XPath Full-Text—Requirements*, 2004, W3C, Working Draft. Available from: <http://www.w3.org/TR/xquery-full-text-requirements/>.
- [63] A. b.v. and S. A. Ltd. *The SeRQL query language*, chapter 5. Aduna b.v., Sirma AI Ltd., 2002. Available from: <http://www.openrdf.org/doc/users/ch05.html>.
- [64] L. Cardelli. *Type systems*, chapter 103, pages 2208-2236. ACM Press, 1996. Available from: <http://research.microsoft.com/Users/luca/Papers/TypeSystems.pdf>.
- [65] J. Carroll, C. Bizer, P. Hayes, and P. Stickler. Named Graphs, Provenance and Trust. Technical Report HPL-2004-57, HP Labs, 2004. Available from: <http://www.hp1.hp.com/techreports/2004/HPL-2004-57>.
- [66] J. Carroll and P. Stickler. TriX: RDF Triples in XML (online). May 2004 (retrieved at 2004-07-09). Available from: <http://www.hp1.hp.com/techreports/2004/HPL-2004-56>.
- [67] S. Ceri, G. Gottlob, and L. Tanca. *Logic Programming and Databases*. Springer Verlag, 1990.

- [68] D. Chamberlin, P. Fankhauser, D. Florescu, M. Marchiori, and J. Robie, editors. *XML Query Use Cases*, 2003, W3C, Working Draft. Available from: <http://www.w3.org/TR/xquery-use-cases/>.
- [69] D. Chamberlin, P. Fankhauser, M. Marchiori, and J. Robie, editors. *XML Query (XQuery) Requirements*, 2003, W3C, Working Draft. Available from: <http://www.w3.org/TR/xquery-requirements/>.
- [70] D. Chamberlin, J. Robie, and D. Florescu. Quilt: An XML Query Language for Heterogeneous Data Sources. In *Proc. Workshop on Web and Databases*, 2000.
- [71] A. Chandra and D. Harel. Computable queries for relational data bases (Preliminary Report). *Journal of Computer Systems and Sciences*, 21:156-178, 1980. Available from: <http://portal.acm.org/citation.cfm?id=804423>.
- [72] A. Chandra and D. Harel. Structure and Complexity of Relational Queries. *Journal of Computer Systems and Sciences*, 25(1):99-128, 1982. Available from: <http://www.wisdom.weizmann.ac.il/~dharel/SCANNED.PAPERS/RelationalQueries.pdf>.
- [73] Y. Chen and P. Revesz. CXQuery: A Novel XML Query Language. In *Proc. International Conference on Advances in Infrastructure for Electronic Business, Science, and Medicine on the Internet*, 2002. Available from: <http://citeseer.ist.psu.edu/rd/22271714%2C539624%2C1%2C0.25%2CDownload/http://citeseer.ist.psu.edu/cache/papers/cs/26473/http:zSzzSzwww.ssgrr.itzSzenzSzssgrr2002wzSzpaperszSz216.pdf/chen02cxquery.pdf>.
- [74] V. Christophides, S. Cluet, and G. Moerkotte. Evaluating Queries with Generalized Path Expressions. In *Proc. ACM SIGMOD International Conference on Management of Data*, pages 413-422, 1996. Available from: <http://citeseer.ist.psu.edu/rd/22271714%2C52273%2C1%2C0.25%2CDownload/http://citeseer.ist.psu.edu/cache/papers/cs/4384/http:zSzzSzwww-rocq.inria.frzSzwhozSzVassilis.ChristophideszSzsigmod96.pdf/christophides96evaluating.pdf>.
- [75] V. Christophides, G. Karvounarakis, I. Koffina, G. Kokkinidis, A. Magkanaraki, D. Plexousakis, G. Serfiotis, and V. Tannen. The ICS-FORTH SWIM: A Powerful Semantic Web Integration Middleware. In *Proc. International Workshop on Semantic Web and Databases*, 2003. Available from: <http://athena.ics.forth.gr:9090/RDF/publications/swdb2003.pdf>.
- [76] V. Christophides, D. Plexousakis, G. Karvounarakis, and S. Alexaki. Declarative Languages for Querying Portal Catalogs. In *Proc. DELOS Workshop: Information Seeking, Searching and Querying in Digital Libraries*, 2000. Available from: <http://athena.ics.forth.gr:9090/RDF/publications/delos2000.pdf>.
- [77] J. Clark, editor. *XSL Transformations (XSLT) Version 1.0*, 1999, W3C, Recommendation. Available from: <http://www.w3.org/TR/xslt/>.
- [78] J. Clark and S. DeRose, editors. *XML Path Language (XPath) Version 1.0*, 1999, W3C, Recommendation. Available from: <http://www.w3.org/TR/xpath/>.

- [79] J. Clark and M. Makoto, editors. *RELAX NG Specification*, 2001, OASIS, Committee Specification. Available from: <http://www.relaxng.org/spec-20011203.html>.
- [80] J. Clark and M. Makoto, editors. *Regular-grammar-based validation—RELAX NG*, 2002, International Organisation for Standardization, Draft International Standard. Available from: http://www.y12.doe.gov/sgml/sc34/document/0362_files/relaxng-is.pdf.
- [81] K. Clark, editor. *RDF Data Access Use Cases and Requirements*, 2004, W3C, Working Draft. Available from: <http://www.w3.org/TR/rdf-dawg-uc/>.
- [82] K. Clark and D. Connolly. *RDF Data Access Design Evaluations* (online). 2004 (retrieved at 2004-08-09). Available from: <http://www.w3.org/2001/sw/DataAccess/DesignEvaluations>.
- [83] E. Codd. A Relational Model of Data for Large Shared Data Banks. *Communications of the ACM*, 13(6):377-387, 1970. Available from: <http://www.acm.org/classics/nov95/toc.html>.
- [84] E. Codd. Relational Completeness of Data Base Sublanguages. In R. Rustin, editor, *Database Systems: Courant Computer Science Symposia 6*. Prentice-Hall, 1972.
- [85] L. Consultants. *Nexus Query Language* (online). 2000 (retrieved at 2004-04-26). Available from: <http://www.langdale.com.au/RDF/NexusQueryLanguage.html>.
- [86] J. Cowan and R. Tobin, editors. *XML Information Set (Second Edition)*, 2004, W3C, Recommendation. Available from: <http://www.w3.org/TR/2004/REC-xml-infoSet-20040204/>.
- [87] R. Daniel, editor. *Harvesting RDF Statements from XLinks*, 2000, W3C, Note. Available from: <http://www.w3.org/TR/xmlink2rdf/>.
- [88] I. Davis. *RDF Template Language 1.0* (online). September 2003 (retrieved at 2004-04-26). Available from: <http://www.semanticplanet.com/2003/08/rdfT/spec>.
- [89] N. Deakin. *ReoPath* (online). 2003 (retrieved at 2004-07-18). Available from: <http://www.xulplanet.com/ndeakin/arts/rpath-fns.txt>.
- [90] S. Decker, D. Brickley, J. Saarela, and J. Angele. A Query and Inference Service for RDF. In *Proc. W3C QL'98 - Query Languages 1998*, December 1998. Available from: <http://www.w3.org/TandS/QL/QL98/pp/queryservice.html>.
- [91] S. DeRose, R. Daniel, E. Maier, and J. Marsh, editors. *XPointer xmlns() Scheme*, 2003, W3C, Recommendation. Available from: <http://www.w3.org/TR/xptr-xmlns/>.
- [92] S. DeRose, E. Maier, and R. Daniel, editors. *XPointer xpointer() Scheme*, 2002, W3C, Working Draft. Available from: <http://www.w3.org/TR/xptr-xpointer/>.
- [93] S. DeRose, E. Maier, and D. Orchard, editors. *XML Linking Language (XLink) Version 1.0*, 2001, W3C, Recommendation. Available from: <http://www.w3.org/TR/xmlink/>.

- [94] A. Deutsch, M. Fernandez, D. Florescu, A. Levy, and D. Suciu. *XML-QL: A Query Language for XML*, 1998, W3C, Note. Available from: <http://www.w3.org/TR/1998/NOTE-xml-ql-19980819/>.
- [95] C. Dollin. Jena Toolkit—Reification Howto (online). 2003 (retrieved at 2004-07-20). Available from: <http://jena.sourceforge.net/how-to/reification.html>.
- [96] E. Dumbill. Putting RDF to Work (online). 2000 (retrieved at 2004-07-17). Available from: <http://www.xml.com/pub/a/2000/08/09/rdfdb/>.
- [97] D. Eastlake and A. Panitz. Reserved Top Level DNS Names. RFC 2606, IETF, 1999. Available from: <http://www.rfc-editor.org/rfc/rfc2606.txt>.
- [98] D. Fallside, editor. *XML Schema Part 0: Primer*, 2001, W3C, Recommendation. Available from: <http://www.w3.org/TR/xmlschema-0/>.
- [99] L. Fegaras, D. Levine, S. Bose, and V. Chaluvadi. Query Processing of Streamed XML Data. In *Proc. International Conference on Information and Knowledge Management*, 2002. Available from: <http://lambda.uta.edu/cikm02.pdf>.
- [100] M. Fernandez, A. Malhotra, J. Marsh, M. Nagy, and N. Walsh, editors. *XQuery 1.0 and XPath 2.0 Data Model*, 2004, W3C, Working Draft. Available from: <http://www.w3.org/TR/xpath-datamodel/>.
- [101] M. Fernandez, J. Simeon, and P. Wadler, editors. *XML Query Languages: Experiences and Exemplars*, 1999, Draft. Available from: <http://www.w3.org/1999/09/ql/docs/xquery.html>.
- [102] R. Fikes, P. Hayes, and I. Horrocks, editors. *DAML Query Language (DQL): Abstract Specification*, 2002, DAML Joint Committee. Available from: <http://www.dam1.org/dql/>.
- [103] R. Fikes, P. Hayes, and I. Horrocks. OWL-QL - A Language for Deductive Query Answering on the Semantic Web. *Journal of Web Semantics*, To appear. Available from: <http://ks1.stanford.edu/projects/owl-ql/>.
- [104] D. Florescu, A. Grünhagen, and D. Kossmann. XL: An XML Programming Language for Web Service Specification and Composition. In *Proc. International World Wide Web Conference*, May 2002. Available from: <http://xl.in.tum.de/pub1/www2002.html>.
- [105] D. Florescu, A. Grünhagen, and D. Kossmann. XL: An XML Programming Language for Web Service Specification and Composition. *Computer Networks*, 42(5), 2003. Available from: <http://xl.informatik.uni-heidelberg.de/pub1/comnet.pdf>.
- [106] D. Florescu, C. Hillery, D. Kossmann, P. Lucas, F. Riccardi, T. Westmann, M. Carey, A. Sundararajan, and G. Agrawal. The BEA/XQRL Streaming XQuery Processor. In *Proc. International Conference on Very Large Databases*, 2003. Available from: <http://www.vldb.org/conf/2003/papers/S30P01.pdf>.
- [107] J. Frohn, G. Lausen, and H. Uphoff. Access to Objects by Path Expressions and Rules. In *Proc. International Conference on Very Large Databases*, 1994.

- [108] N. Fuchs, U. Schwertel, and R. Schwitter. Attempto Controlled English (ACE) Language Manual, Version 3.0. Technical Report 99.03, Department of Computer Science, University of Zurich, 1999. Available from: <http://www.ifi.unizh.ch/attempto/publications/papers/ace3manual.pdf>.
- [109] N. Fuhr and K. Gross. XIRQL: a Query Language for Information Retrieval in XML Documents. In *Proc. ACM Conference on Research and Development in Information Retrieval*, 2001. Available from: <http://portal.acm.org/citation.cfm?id=383985&dl=ACM&coll=portal>.
- [110] L. Garshol. The Linear Topic Map Notation (online, retrieved at 2007-07-18). Available from: <http://www.ontopia.net/download/ltn.html>.
- [111] L. Garshol. tolog—A topic map query language. In *Proc. XML Europe*, 2001. Available from: <http://www.ontopia.net/topicmaps/materials/tolog.html>.
- [112] L. Garshol. Extending tolog—Proposal for tolog 1.0. In *Proc. Extreme Markup Languages*, 2003. Available from: <http://www.ontopia.net/topicmaps/materials/extending-tolog.html>.
- [113] L. Garshol. tolog 0.1. Technical report, Ontopia, 2003. Available from: <http://www.ontopia.net/topicmaps/materials/tolog-spec.html>.
- [114] L. Garshol. tolog—Language tutorial (online). 2004 (retrieved at 2004-04-27). Available from: <http://www.ontopia.net/omnigator/docs/query/tutorial.html>.
- [115] L. Garshol and R. Barta, editors. *Topic Map Query Language, Requirements*, November 2003, ISO/IEC, Draft. Available from: <http://www.y12.doe.gov/sgml/sc34/document/0448.htm>.
- [116] L. M. Garshol. Living with Topic Maps and RDF (online). 2003 (retrieved at 2004-08-19). Available from: <http://www.ontopia.net/topicmaps/materials/tmrdf.html>.
- [117] B. Glimm. A Query Language for Web Ontologies. Bachelor thesis, Hamburg University of Applied Sciences, University of Manchester, 2004. Available from: <http://www2.cs.man.ac.uk/~glimmbx/download/report.pdf>.
- [118] R. Goldman, S. Chawathe, A. Crespo, and J. McHugh. A Standard Textual Interchange Format for the Object Exchange Model (OEM). Technical report, Database Group, Stanford University, 1996. Available from: <http://www-db.stanford.edu/~mchughj/oemsyntax/oemsyntax.html>.
- [119] G. Gottlob, C. Koch, and R. Pichler. Efficient Algorithms for Processing XPath Queries. In *Proc. International Conference on Very Large Databases*, 2002. Available from: <http://www.dbai.tuwien.ac.at/staff/koch/download/vldb2002.pdf>.
- [120] G. Gottlob, C. Koch, and R. Pichler. The Complexity of XPath Query Evaluation. In *Proc. ACM Symposium on Principles of Database Systems*, 2003. Available from: <http://www.dbai.tuwien.ac.at/staff/koch/download/pods2003.pdf>.

- [121] G. Gottlob, C. Koch, and R. Pichler. XPath Processing in a Nutshell. *SIGMOD Record*, 32(2), 2003. Available from: <http://www.dbai.tuwien.ac.at/staff/koch/download/xpnut.pdf>.
- [122] G. Gottlob, C. Koch, and R. Pichler. XPath Query Evaluation: Improving Time and Space Efficiency. In *Proc. International Conference on Data Engineering*, 2003. Available from: <http://www.dbai.tuwien.ac.at/staff/koch/download/icde2003.pdf>.
- [123] J. Grant and D. Beckett, editors. *RDF Test Cases*, February 2004, W3C. Available from: <http://www.w3.org/TR/rdf-testcases/>.
- [124] P. Grosso, E. Maier, J. Marsh, and N. Walsh, editors. *XPointer element() Scheme*, 2003, W3C, Recommendation. Available from: <http://www.w3.org/TR/xptr-element/>.
- [125] P. Grosso, E. Maier, J. Marsh, and N. Walsh, editors. *XPointer Framework*, 2003, W3C, Recommendation. Available from: <http://www.w3.org/TR/xptr-framework/>.
- [126] H. L. S. W. R. Group. Jena - A Semantic Web Framework for Java (online, retrieved at 2004-07-08). Available from: <http://jena.sourceforge.net/>.
- [127] R. Guha. *Contexts: A Formalization and Some Applications*. PhD thesis, Stanford University, 1995.
- [128] R. Guha. rdfDB Query Language (online). 2000 (retrieved at 2004-04-26). Available from: <http://www.guha.com/rdfdb/query.html>.
- [129] R. Guha, O. Lassila, E. Miller, and D. Brickley. Enabling Inferencing. In *Proc. W3C QL'98 - Query Languages 1998*, December 1998.
- [130] J. Gylta. XTMPath, Manipulating Topic Map Data Structures (online). 2002 (retrieved at 2004-08-19). Available from: <http://topicmaps.it.bond.edu.au/docs/13>.
- [131] P. Haase, J. Broekstra, A. Eberhart, and R. Volz. A Comparison of RDF Query Languages. In *Proc. International Semantic Web Conference*, 2004. Available from: <http://www.aifb.uni-karlsruhe.de/WBS/pha/rdf-query/rdfquery.pdf>.
- [132] M. Harren, M. Raghavachari, O. Shmueli, M. Burke, V. Sarkar, and R. Bordawekar. XJ: Integration of XML Processing into Java. In *Proc. International World Wide Web Conference*, 2004. Available from: <http://www2004.org/proceedings/docs/2p340.pdf>.
- [133] S. Harris and N. Gibbins. 3store: Efficient Bulk RDF Storage. In *Proc. International Workshop on Practical and Scalable Semantic Systems*, 2003. Available from: <http://km.aifb.uni-karlsruhe.de/ws/psss03/proceedings/harris-et-al.pdf>.
- [134] A. Harth. Triple Tutorial (online). 2004 (retrieved at 2004-08-01). Available from: <http://triple.semanticweb.org/doc/tutorial.html>.
- [135] P. Hayes and B. McBride, editors. *RDF Semantics*, 2004, W3C, Recommendation. Available from: <http://www.w3.org/TR/rdf-mt/>.

- [136] I. Horrocks and P. Patel-Schneider. Reducing OWL Entailment to Description Logic Satisfiability. In *Proc. International Semantic Web Conference*, pages 17–29. Springer-Verlag, 2003. Available from: <http://www.cs.man.ac.uk/~horrocks/ISWC2003/HoPa03b.html>.
- [137] I. Horrocks and P. Patel-Schneider. A Proposal for an OWL Rules Language. In *Proc. International World Wide Web Conference*, pages 723–731. ACM, 2004. Available from: <http://www.cs.man.ac.uk/~horrocks/DAML/Rules/>.
- [138] I. Horrocks, P. Patel-Schneider, H. Boley, S. Tabet, B. Grosf, and M. Dean. *SWRL: A Semantic Web Rule Language—Combining OWL and RuleML*, 2004, W3C, Member submission. Available from: <http://www.w3.org/Submission/2004/SUBM-SWRL-20040521/>.
- [139] I. Horrocks and S. Tessaris. A Conjunctive Query Language for Description Logic ABoxes. In *Proc. National Conference on Artificial Intelligence*, 2000. Available from: <http://www.cs.man.ac.uk/~tessaris/papers/aaai00.pdf>.
- [140] I. Horrocks, F. van Harmelen, and P. Patel-Schneider, editors. *DAML+OIL*, 2001, Joint US/EU ad hoc Agent Markup Language Committee, Revised Language Specification. Available from: <http://www.daml.org/2001/03/daml+oil-index.html>.
- [141] H. Hosoya and B. Pierce. XDuce: A Typed XML Processing Language. *ACM Transactions on Internet Technology*, 3(2):117–148, 2003.
- [142] J. Hynynen and O. Lassila. On the Use of Object-Oriented Paradigm in a Distributed Problem Solver. *AI Communications*, 2(3):142–151, 1989.
- [143] T. D. C. M. Initiative. DCMI term declarations represented in RDF schema language (online). March 2003 (retrieved at 2004-07-09). Available from: <http://dublincore.org/schemas/rdfs/>.
- [144] *ISO/IEC 13250 Topic Maps*, International Organization for Standardization, International Standard. Available from: http://www.y12.doe.gov/sgml/sc34/document/0322_files/iso13250-2nd-ed-v2.pdf.
- [145] G. Karvounarakis, S. Alexaki, V. Christophides, D. Plexousakis, and M. Scholl. RQL: A Declarative Query Language for RDF. In *Proc. International World Wide Web Conference*, May 2002. Available from: <http://athena.ics.forth.gr:9090/RDF/publications/www2002/www2002.html>.
- [146] G. Karvounarakis, V. Christophides, D. Plexousakis, and S. Alexaki. Querying RDF Descriptions for Community Web Portals. In *Proc. Journees Bases de Donnees Avancees*, 2001. Available from: <http://athena.ics.forth.gr:9090/RDF/publications/bda2001.pdf>.
- [147] G. Karvounarakis, A. Magkanaraki, S. Alexaki, V. Christophides, D. Plexousakis, M. Scholl, and K. Tolle. Querying the Semantic Web with RQL. *Computer Networks and ISDN Systems Journal*, 42(5):617–640, August 2003. Available from: <http://athena.ics.forth.gr:9090/RDF/publications/comp-networks2003.pdf>.

- [148] G. Karvounarakis, A. Magkanaraki, S. Alexaki, V. Christophides, D. Plexousakis, M. Scholl, and K. Tolle. RQL: A Functional Query Language for RDF. In P. Gray, P. King, and A. Poulouvasilis, editors, *The Functional Approach to Data Management*, chapter 18, pages 435–465. Springer-Verlag, 2004. Available from: <http://athena.ics.forth.gr:9090/RDF/publications/FuncBook.pdf>.
- [149] H. Katz. XsRQL: an XQuery-style Query Language for RDF (online). 2004 (retrieved at 2004-07-21). Available from: <http://www.fatdog.com/xsrql.html>.
- [150] M. Kay, editor. *XSL Transformations (XSLT) Version 2.0*, 2003, W3C, Working Draft. Available from: <http://www.w3.org/TR/xslt20/>.
- [151] S. Kepser. A Simple Proof of the Turing-Completeness of XSLT and XQuery. In *Proc. Extreme Markup Languages*, 2004. Available from: <http://tcl.sfs.uni-tuebingen.de/~kepsker/papers/xsltxq.pdf>.
- [152] M. Kifer, G. Lausen, and J. Wu. Logical Foundations of Object Oriented and Frame Based Languages. *Journal of ACM*, 42:741–843, 1995.
- [153] C. Kirchner, Z. Oian, P. Singh, and J. Stuber. Xemantics: a Rewriting Calculus-Based Semantics of XSLT. Technical Report A01-R-386, LORIA, 2001. Available from: <http://www.loria.fr/~stuber/publications/KirchnerQianSinghStuber2001LORIA.html>.
- [154] G. Klyne. Contexts for RDF Information Modelling (online). 2000 (retrieved at 2004-07-16). Available from: <http://www.ninebynine.org/RDFNotes/RDFContexts.html>.
- [155] G. Klyne and J. Carroll, editors. *Resource Description Framework (RDF): Concepts and Abstract Syntax*, February 2004, W3C. Available from: <http://www.w3.org/TR/2004/REC-rdf-concepts-20040210/>.
- [156] G. Klyne, J. Carroll, and B. McBride, editors. *Resource Description Framework (RDF): Concepts and Abstract Syntax*, 2004, W3C, Recommendation. Available from: <http://www.w3.org/TR/rdf-concepts/>.
- [157] C. Koch, S. Scherzinger, N. Schweikardt, and B. Stegmaier. FluXQuery: An Optimizing XQuery Processor for Streaming XML Data. In *Proc. International Conference on Very Large Databases*, 2004. Available from: http://www.dbai.tuwien.ac.at/staff/koch/download/vldb2004_demo.pdf.
- [158] R. Ksiezzyk. Answer is just a question [of matching Topic Maps]. In *Proc. XML Europe*, 2000. Available from: <http://www.gca.org/papers/xmleurope2000/papers/s22-03.html>.
- [159] M. Lacher and S. Decker. On the Integration of Topic Maps and RDF Data. In *Proc. Extreme Markup Languages*, 2001. Available from: <http://www.semanticweb.org/SWWS/program/full/paper53.pdf>.
- [160] M. Lacher and S. Decker. RDF, Topic Maps, and the Semantic Web. *Markup Languages: Theory and Practice*, 3(3):313–331, December 2001. Available from: <https://206.191.28.118/docushare/dsweb/GetRepr/Document-1157/html>.

- [161] O. Lassila. BEEF Reference Manual—A Programmer’s Guide to the BEEF Frame System, Second Version. Technical Report HTKK-TKO-C46, Department of Computer Science, Helsinki University of Technology, 1991.
- [162] O. Lassila. Enabling Semantic Web Programming by Integrating RDF and Common Lisp. In *Proc. Semantic Web Working Symposium*, July 2001. Available from: <http://www.lassila.org/publications/2001/swws-01-abstract.shtml>.
- [163] O. Lassila. Taking the RDF Model Theory Out for a Spin. In *Proc. Semantic Web Working Symposium*, June 2002. Available from: <http://www.lassila.org/publications/2002/lassila-iswc2002-abstract.shtml>.
- [164] O. Lassila. Ivanhoe: an RDF-Based Frame System (online). 2004 (retrieved at 2004-07-17). Available from: <http://wilbur-rdf.sourceforge.net/docs/ivanhoe.html>.
- [165] O. Lassila. Wilbur Query Language Comparison (online). 2004 (retrieved at 2004-07-17). Available from: <http://wilbur-rdf.sourceforge.net/2004/05/11-comparison.shtml>.
- [166] O. Lassila. Wilbur Semantic Web Toolkit (online). 2004 (retrieved at 2004-07-17). Available from: <http://wilbur-rdf.sourceforge.net/docs/>.
- [167] O. Lassila and R. Swick, editors. *Resource Description Framework (RDF) Model and Syntax Specification*, 1999, W3C, Recommendation. Available from: <http://www.w3.org/TR/1999/REC-rdf-syntax-19990222/>.
- [168] A. Laux and L. Martin. *XUpdate—XML Update Language*, 2000, XML:DB Initiative, Working Draft. Available from: <http://xmldb-org.sourceforge.net/xupdate/xupdate-wd.html>.
- [169] E. Lenz. What’s New in XSLT 2.0 (online). 2002 (retrieved at 2004-07-08). Available from: <http://www.xml.com/pub/a/2002/04/10/xslt2.html>.
- [170] M. Liu. A Logical Foundation for XML. In *Proc. International Conference on Advanced Information Systems Engineering*. Springer-Verlag, 2002. Available from: <http://springerlink.metapress.com/app/home/contribution.asp?wasp=2pw78ah4fg7xrjbd41vm&referrer=parent&backto=issue,39,69;journal,764,1660;linkingpublicationresults,1:105633,1>.
- [171] B. Ludäscher, R. Himmeroeder, G. Lausen, W. May, and C. Schleppehorst. Managing Semistructured Data with FLORID: A Deductive Object-oriented Perspective. *Information Systems*, 23(8):1-25, 1998.
- [172] B. Ludäscher, Y. Papakonstantinou, and P. Velikhov. *A Brief Introduction to XMAS*, 1999, Database Group, University of California, San Diego. Available from: <http://www.db.ucsd.edu/Projects/MIX/docs/XMAS-intro.pdf>.
- [173] R. Luk, H. Leong, T. Dillon, A. Chan, W. B. Croft, and J. Allan. A Survey in Indexing and Searching XML Documents. *Journal of the American Society for Information Science and Technology*, 53(6):415-437, 2002.

- [174] A. Magkanaraki, G. Karvounarakis, V. Christophides, D. Plexousakis, and T. Anh. Ontology Storage and Querying. Technical Report 308, Foundation for Research and Technology Hellas, April 2002. Available from: <http://139.91.183.30:9090/RDF/publications/tr308.pdf>.
- [175] A. Magkanaraki, V. Tannen, V. Christophides, and D. Plexousakis. Viewing the Semantic Web Through RVL Lenses. In *Proc. International Semantic Web Conference*, October 2003. Available from: <http://athena.ics.forth.gr:9090/RDF/publications/iswc03.pdf>.
- [176] D. Maier. Database Desiderata for an XML Query Language. In *Proc. W3C QL'98 - Query Languages 1998*, December 1998. Available from: <http://www.w3.org/TandS/QL/QL98/pp/maier.html>.
- [177] A. Malhotra, J. Melton, J. Robie, and M. Rys, editors. *XML Syntax for XQuery 1.0 (XQueryX)*, 2003, W3C, Working Draft. Available from: <http://www.w3.org/TR/xqueryx/>.
- [178] A. Malhotra, J. Melton, and N. Walsh, editors. *XQuery 1.0 and XPath 2.0 Functions and Operators*, 2004, W3C, Working Draft. Available from: <http://www.w3.org/TR/xpath-functions/>.
- [179] F. Manola, E. Miller, and B. McBride, editors. *RDF Primer*, 2004, W3C, Recommendation. Available from: <http://www.w3.org/TR/rdf-primer/>.
- [180] M. Marchiori. The Pseudo Natural Language Interface (online, retrieved at 2004-07-17). Available from: <http://www.w3.org/RDF/MetaLog/docs/pnl.html>.
- [181] M. Marchiori, A. Epifani, and S. Trevisan. Metalog v2.0: Quick User Guide. Technical report, W3C, 2004. Available from: <http://www.w3.org/RDF/MetaLog/>.
- [182] M. Marchiori and J. Saarela. Query + Metadata + Logic = Metalog. In *Proc. W3C QL'98 - Query Languages 1998*, December 1998. Available from: <http://www.w3.org/TandS/QL/QL98/pp/metalog.html>.
- [183] M. Marchiori and J. Saarela. Towards the Semantic Web: Metalog (online). 1999 (retrieved at 2004-04-26). Available from: <http://www.w3.org/RDF/MetaLog/CIKM-050299.html>.
- [184] J. Marsh and D. Orchard, editors. *XML Inclusions (XInclude) Version 1.0*, 2004, W3C, Candidate Recommendation. Available from: <http://www.w3.org/TR/xinclude/>.
- [185] K. Matsuyama, M. Kraus, K. Kitagawa, and N. Saito. A Path-Based RDF Query Language for CC/PP and UAProf. In *Proc. IEEE Conference on Pervasive Computing and Communications Workshops*, 2004.
- [186] W. May. LoPiX: a System for XML data Integration and Manipulation. In *Proc. International Conference on Very Large Databases*, 2001.
- [187] W. May. XPath-Logic and XPathLog: A Logic-Programming Style XML Data Manipulation Language. *Theory and Practice of Logic Programming*, 3(4):499-526, 2004. Available from: <http://dbis.informatik.uni-goettingen.de/Publics/04/TPLP04.html>.

- [188] D. McGuinness and F. van Harmelen, editors. *OWL Web Ontology Language—Overview*, 2004, W3C, Recommendation. Available from: <http://www.w3.org/TR/owl-features/>.
- [189] E. Meijer, W. Schulte, and G. Bierman. Programming with Circles, Triangles and Rectangles. In *Proc. XML Conference and Exhibition*, 2003.
- [190] E. Meijer and M. Shields. XMLambda: A functional language for constructing and manipulating XML documents (online). 1999 (retrieved at 2004-08-09). Available from: <http://www.cse.ogi.edu/~mbs/pub/xmlambda/>.
- [191] S. Melnik. Simplified Syntax for RDF (online). December 1999 (retrieved at 2004-07-09). Available from: <http://www-db.stanford.edu/~melnik/rdf/syntax.html>.
- [192] S. Melnik. Representing UML in RDF (online). 2000 (retrieved at 2004-08-19). Available from: <http://www-db.stanford.edu/~melnik/rdf/uml/>.
- [193] L. Miller. Inkling: RDF query using SquishQL (online, retrieved at 2004-07-08). Available from: <http://sw1.illrt.org/rdfquery/>.
- [194] L. Miller, A. Seaborne, and A. Reggiori. Three Implementations of SquishQL, a Simple RDF Query Language. In *Proc. International Semantic Web Conference*, June 2002. Available from: <http://rdfstore.sourceforge.net/documentation/papers/HPL-2002-110.pdf>.
- [195] W. Nejdl, B. Wolf, C. Qu, S. Decker, M. Sintek, A. Naeve, M. Nilsson, M. Palmer, and T. Risch. Edutella: A P2P Networking Infrastructure Based on RDF. In *Proc. International World Wide Web Conference*, May 2002. Available from: <http://www.kbs.uni-hannover.de/Arbeiten/Publikationen/2002/p597-nejdl.pdf>.
- [196] M. Nilsson and W. Siberski. RDF Query Exchange Language (QEL): Concepts, Semantics and RDF Syntax (online). February 2004 (retrieved at 2004-04-26). Available from: <http://edutella.jxta.org/spec/qel.html>.
- [197] M. Nilsson, W. Siberski, and J. Tane. Edutella Retrieval Service: Concepts and RDF Syntax (online). June 2004 (retrieved at 2004-04-26). Available from: <http://edutella.jxta.org/spec/retrieval.html>.
- [198] A. Novoselsky and K. Karun. XSLTVM—an XSLT Virtual Machine. In *Proc. XML Europe*, 2000. Available from: <http://www.gca.org/papers/xml europe2000/papers/s35-03.html>.
- [199] M. Odersky. Report on the Programming Language Scala. Technical report, Ecole Polytechnique Federale de Lausanne, 2002. Available from: <http://lamp.epfl.ch/~odersky/scala/>.
- [200] U. Ogbuji. Versa by example (online, retrieved at 2004-04-26). Available from: <http://uche.ogbuji.net/tech/rdf/versa/versa-by-example.txt>.
- [201] U. Ogbuji. Thinking XML: Basic XML and RDF techniques for knowledge management: Part 6: RDF Query using Versa (online). April 2002 (retrieved at 2004-04-26). Available from: <http://www-106.ibm.com/developerworks/xml/library/x-think10/>.

- [202] R. Oldakowski and C. Bizer. RAP: RDF API for PHP. In *Proc. International Workshop on Interpreted Languages*, 2004. Available from: <http://www.wiwiss.fu-berlin.de/suhl/radek/pub/rap-oldakowski.pdf>.
- [203] F. Olken and J. McCarthy. Requirements and Desiderata for an XML Query Language. In *Proc. W3C QL'98 - Query Languages 1998*, December 1998.
- [204] M. Olson and U. Ogbuji. Versa Specification (online). 2003 (retrieved at 2004-04-26). Available from: http://uche.ogbuji.net/tech/rdf/versa/versa.doc?xslt=/ftss/data/docbook_html1.xslt.
- [205] D. Olteanu, T. Furche, and F. Bry. An Efficient Single-Pass Query Evaluator for XML Data Streams. In *Proc. ACM Symposium on Applied Computing*, 2004. Available from: <http://www.pms.ifi.lmu.de/publikationen#PMS-FB-2004-1>.
- [206] D. Olteanu, T. Furche, and F. Bry. Evaluating Complex Queries against XML streams with Polynomial Combined Complexity. In *Proc. British National Conference on Databases*, 2004. Available from: <http://www.pms.ifi.lmu.de/publikationen#PMS-FB-2003-15>.
- [207] D. Olteanu, H. Meuss, T. Furche, and F. Bry. XPath: Looking Forward. In *Proc. EDBT Workshop on XML Data Management*, LNCS 2490. Springer Verlag, 2002. Available from: <http://www.pms.ifi.lmu.de/publikationen/#PMS-FB-2002-4>.
- [208] S. Palmer. Pondering RDF Path (online). 2003 (retrieved at 2004-04-26). Available from: <http://infomesh.net/2003/rdfpath/>.
- [209] Y. Papakonstantinou, H. Garcia-Molina, and J. Widom. Object Exchange across Heterogeneous Information Sources. In *Proc. International Conference on Data Engineering*, pages 251-260, 1995.
- [210] P. Patel-Schneider, P. Hayes, and I. Horrocks, editors. *OWL Web Ontology Language—Semantics and Abstract Syntax*, 2004, W3C, Recommendation. Available from: <http://www.w3.org/TR/owl-semantics/>.
- [211] P. Patel-Schneider and J. Simeon. The Yin/Yang Web: XML Syntax and RDF Semantics. In *Proc. International World Wide Web Conference*, May 2002.
- [212] S. Pepper and G. Moore, editors. *XML Topic Maps (XTM) 1.0*, 2001, TopicMaps.org, Specification. Available from: <http://www.topicmaps.org/xtm/index.html>.
- [213] R. Pinchuk. Toma - Topic Map Query Language (online). 2004 (retrieved at 2004-04-27). Available from: <http://www.spaceapplications.com/toma/>.
- [214] M. Plusch. *Water: Simplified Web Services and XML Programming*. Wiley, 2002. Available from: http://waterlanguage.org/water_book_2002/.
- [215] E. Prud'hommeaux. Algae2 HOTWO (online). 2003 (retrieved at 2004-04-26). Available from: <http://www.w3.org/2004/05/06-Algae/>.
- [216] E. Prud'hommeaux. Algae Extension for Rules (online). 2004 (retrieved at 2004-07-18). Available from: <http://www.w3.org/2004/06/20-rules/>.

- [217] E. Prud'hommeaux. Algae RDF Query Language (online). 2004 (retrieved at 2004-07-18). Available from: <http://www.w3.org/2004/05/06-Algae/>.
- [218] E. Prud'hommeaux. Optimized RDF Access to Relational Databases (online). 2004 (retrieved at 2004-07-18). Available from: <http://www.w3.org/2004/04/30-RDF-RDB-access/>.
- [219] E. Prud'hommeaux and B. Grosf. RDF Query Survey (online). 2004 (retrieved at 2004-04-26). Available from: <http://www.w3.org/2001/11/13-RDF-Query-Rules/>.
- [220] E. Prud'hommeaux and A. Seaborne. BRQL - A Query Language for RDF (online, retrieved at 2004-07-20). Available from: <http://www.w3.org/2004/07/08-BRQL/>.
- [221] A. Reggiori and A. Seaborne. RDF Query and Rule languages: Use Cases and Examples (online). 2004 (retrieved at 2004-04-26). Available from: <http://rdfstore.sourceforge.net/2002/06/24/rdf-query/query-use-cases.html>.
- [222] A. Reggiori and D.-W. van Gulik. RDFStore—Perl API for RDF Storage (online, retrieved at 2004-07-18). Available from: <http://rdfstore.sourceforge.net/>.
- [223] D. Reynolds. RDF-QBE: a Semantic Web Building Block. Technical Report HPL-2002-327, HP Labs, 2002. Available from: <http://www.hp1.hp.com/techreports/2002/HPL-2002-327.html>.
- [224] J. Robie. The Syntactic Web: Syntax and Semantics on the Web. In *Proc. XML Conference and Exposition*, December 2001. Available from: <http://www.idealliance.org/papers/xml2001/papers/html/03-01-04.html>.
- [225] J. Robie, E. Derksen, P. Frankhauser, E. Howland, G. Huck, I. Macherius, M. Murata, M. Resnick, and H. Schöning. XQL (XML Query Language) (online). 1999 (retrieved at 2004-06-05). Available from: <http://www.ibiblio.org/xql/xql-proposal.html>.
- [226] J. Robie, L. M. Garshol, S. Newcomb, M. Fuchs, L. Miller, D. Brickley, V. Christophides, and G. Karvounarakis. The Syntactic Web: Syntax and Semantics on the Web. *Markup Languages: Theory and Practice*, 3(4):411-440, 2001. Available from: <http://www.w3.org/XML/2002/08/robie.syntacticweb.html>.
- [227] J. Robie, J. Lapp, and D. Schach. XML Query Language (XQL). In *Proc. W3C QL'98 - Query Languages 1998*, December 1998. Available from: <http://www.w3.org/TandS/QL/QL98/pp/xql.html>.
- [228] P. V. Roy and S. Haridi. *Concepts, Techniques, and Models of Computer Programming*. MIT Press, 2004.
- [229] S. Russell and S. Palmer. Quads (online). 2002 (retrieved at 2004-07-20). Available from: <http://robustai.net/sailor/grammar/Quads.html>.
- [230] S. Schaffert and F. Bry. Querying the Web Reconsidered: A Practical Introduction to Xcerpt. In *Proc. Extreme Markup Languages*, August 2004. Available from: <http://www.pms.informatik.uni-muenchen.de/publikationen/PMS-FB/PMS-FB-2004-7.pdf>.

- [231] A. Seaborne. RDQL - RDF Data Query Language (online, retrieved at 2004-07-08). Available from: <http://www.hp1.hp.com/semweb/rdql.htm>.
- [232] A. Seaborne. RDQL - A Query Language for RDF (online). January 2004 (retrieved at 2004-04-26). Available from: <http://www.w3.org/Submission/2004/SUBM-RDQL-20040109/>.
- [233] A. Seaborne. A Programmer's Introduction to RDQL (online). 2002 April (retrieved at 2004-04-26). Available from: <http://www.hp1.hp.com/semweb/doc/tutorial/RDQL/>.
- [234] R. Shearer. REX evaluation (online, retrieved at 2004-07-18). Available from: <http://lists.w3.org/Archives/Public/public-rdf-dawg/2004AprJun/0572.html>.
- [235] M. Sheshagiri and A. Kunjithapatham. A FIPA Compliant Query Mechanism Using DAML Query Language (DQL) (online, retrieved at 2004-07-27). Available from: <http://www.cs.umbc.edu/~finin/papers/dqlFIPA.html>.
- [236] M. Sintek and S. Decker. TRIPLE—An RDF Query, Inference, and Transformation Language. In *Proc. Deductive Database and Knowledge Management*, October 2001. Available from: <http://triple.semanticweb.org/>.
- [237] M. Sintek and S. Decker. TRIPLE—A Query, Inference, and Transformation Language for the Semantic Web. In *Proc. International Semantic Web Conference*, June 2002. Available from: <http://triple.semanticweb.org/>.
- [238] M. Smith, C. Welty, and D. McGuinness, editors. *OWL Web Ontology Language—Guide*, 2004, W3C, Recommendation. Available from: <http://www.w3.org/TR/owl-guide/>.
- [239] A. Souzis. RxML 1.0 Specification (online). 2004 (retrieved at 2004-04-17). Available from: <http://rx4rdf.liminalzone.org/RxMLSpecification>.
- [240] A. Souzis. RxPath (online). 2004 (retrieved at 2004-04-17). Available from: <http://rx4rdf.liminalzone.org/RxPath>.
- [241] A. Souzis. RxPath Specification Proposal (online). 2004 (retrieved at 2004-04-17). Available from: <http://rx4rdf.liminalzone.org/RxPathSpec>.
- [242] A. Souzis. RxSLT (online). 2004 (retrieved at 2004-04-17). Available from: <http://rx4rdf.liminalzone.org/RxSLT>.
- [243] A. Souzis. RxUpdate (online). 2004 (retrieved at 2004-04-17). Available from: <http://rx4rdf.liminalzone.org/RxUpdate>.
- [244] M. Sperberg-McQueen. How can Tom butter his bread with a knife, if there is a dearth of bread in the larder?, 2001. Available from: <http://xml.coverpages.org/RDF-TopicMaps-LateLazyVersusEarlyPreemptiveReification.html>.
- [245] D. Steer. TreeHugger 1.0 Introduction (online). 2003 (retrieved at 2004-04-26). Available from: <http://www.semanticplanet.com/2003/08/rdft/spec>.
- [246] P. Stickler. CBD—Concise Bounded Description (online). 2004 (retrieved at 2004-08-20). Available from: <http://swdev.nokia.com/uriqa/CBD.html>.

- [247] A. Swartz. RDFPath Proposal (online). 2001 (retrieved at 2004-07-18). Available from: <http://logicerror.com/RDFPathProposal>.
- [248] I. Tatarinov, Z. Ives, A. Halevy, and D. Weld. Updating XML. In *Proc. ACM SIGMOD*, 2001. Available from: <http://www.cis.upenn.edu/~zives/research/UpdatingXML.pdf>.
- [249] S. Tessaris. *Questions and Answers: Reasoning and Querying in Description Logic*. PhD thesis, University of Manchester, Department of Computer Science, April 2001.
- [250] H. Thompson, D. Beech, M. Maloney, and N. Mendelsohn, editors. *XML Schema Part 1: Structures*, 2001, W3C, Recommendation. Available from: <http://www.w3.org/TR/xmlschema-1/>.
- [251] K. Tolle and F. Wleklinski. easy RDF Query Language (eRQL) (online, retrieved at 2004-07-08). Available from: <http://www.dbis.informatik.uni-frankfurt.de/~tolle/RDF/eRQL/>.
- [252] J. Ullman. *Principles of Database and Knowledge-Base Systems*. W. H. Freeman, 1990.
- [253] P. Wadler. A formal semantics of patterns in XSLT (online). 2000 (retrieved at 2004-07-18). Available from: <http://homepages.inf.ed.ac.uk/wadler/papers/xsl-semantics/xsl-semantics.pdf>.
- [254] P. Wadler. Two semantics for XPath (online). 2000 (retrieved at 2004-07-18). Available from: <http://homepages.inf.ed.ac.uk/wadler/papers/xpath-semantics/xpath-semantics.pdf>.
- [255] G. Wagner. Seven Golden Rules for a Web Rule Language. *IEEE Intelligent Systems*, 18(5), 2003.
- [256] M. Wallace and C. Runciman. Haskell and XML: Generic Combinators or Type-Based Translation. In *Proc. International Conference on Functional Programming*, 1999. Available from: <http://www.cs.york.ac.uk/fp/HaXml/icfp99.html>.
- [257] N. Walsh. RDF Twig: accessing RDF graphs in XSLT. In *Proc. Extreme Markup Languages*, 2003. Available from: <http://www.mulberrytech.com/Extreme/Proceedings/xslfo-pdf/2003/Walsh01/EML2003Walsh01.pdf>.
- [258] U. Wiger. XMErl—Interfacing XML and Erlang. In *Proc. International Erlang User Conference*, 2000.
- [259] C. Wilper. RIDIQL Reference (online). 2004 (retrieved at 2004-04-17). Available from: <http://rdq1plus.sourceforge.net/doc/ridiql.html>.
- [260] M. Zoof. Query By Example. In *Proc. AFIPS National Computer Conference*, 1975.
- [261] M. Zoof. Query By Example: A Data Base Language. *IBM Systems Journal*, 16(4):324-343, 1977. Available from: <http://www.research.ibm.com/journal/sj/164/ibmsj1604C.pdf>.