Translation of Heterogeneous Databases into RDF, and Application to the Construction of a SKOS Taxonomical Reference

Franck Michel, Loïc Djimenou, Catherine Faron Zucker, Johan Montagnat

► To cite this version:

Franck Michel, Loïc Djimenou, Catherine Faron Zucker, Johan Montagnat. Translation of Heterogeneous Databases into RDF, and Application to the Construction of a SKOS Taxonomical Reference. Web Information Systems and Technologies: 11th International Conference, WEBIST 2015, Revised Selected Papers, 246, Springer, pp.275-296, 2016, Lecture Notes in Business Information Processing, 978-3-319-30995-8. 10.1007/978-3-319-30995-5_14 . hal-01337308

HAL Id: hal-01337308

<https://hal.science/hal-01337308>

Submitted on 24 Jun 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Translation of Heterogeneous Databases into RDF, and Application to the Construction of a SKOS Taxonomical Reference

Franck Michel¹, Loïc Djimenou¹, Catherine Faron-Zucker¹, and Johan Montagnat¹

Univ. Nice Sophia Antipolis, CNRS, I3S (UMR 7271), France

Abstract. While the data deluge accelerates, most of the data produced remains locked in deep Web databases. For the linked open data to benefit from the potential represented by this huge amount of data, it is crucial to come up with solutions to expose heterogeneous databases as linked data. The xR2RML mapping language is an endeavor towards this goal: it is designed to map various types of databases to RDF, by flexibly adapting to heterogeneous query languages and data models while remaining free from any specific language. It extends R2RML, the W3C recommendation for the mapping of relational databases to RDF, and relies on RML for the handling of various data formats.

In this paper we present xR2RML, we analyse data models of several modern databases as well as the format in which query results are returned, and we show how xR2RML translates any result data element into RDF, relying on existing languages such as XPath and JSONPath when necessary. We illustrate some features of xR2RML such as the generation of RDF collections and containers, and the ability to deal with mixed data formats. We also describe a real-world use case in which we applied xR2RML to build a SKOS thesaurus aimed at supporting studies on History of Zoology, Archaeozoology and Conservation Biology.

Keywords: Linked Data, RDF, R2RML, NoSQL, History of Zoology

1 Introduction

The web of data is now emerging through the publication and interlinking of various open data sets in RDF. Initiatives such as the W3C Data Activity¹ and the Linking Open Data (LOD) project² aim at Web-scale data integration and processing, assuming that making heterogeneous data available in a common machine-readable format should create opportunities for novel applications and services. Their success largely depends on the ability to reach data from the deep web [1], a part of the web content consisting of documents and databases hardly linked with other data sources and hardly indexed by standard

¹ <http://www.w3.org/2013/data/>

² <http://linkeddata.org/>

search engines. The deep web keeps growing as data is continuously accumulated in ever more heterogeneous databases. In particular, NoSQL systems have gained a remarkable success during recent years. Driven by major web companies, they have been developed to meet requirements of web 2.0 services, that relational databases (RDB) could not achieve (flexible schema, high throughput, high availability, horizontal elasticity on commodity hardware). Thus, NoSQL systems should now be considered as potential heavy contributors of linked open data. Other types of databases have been developed over time, either for generic purpose or specific domains, such as XML databases (notably used in edition and digital humanities), object-oriented databases or directory-based databases.

Significant efforts have been invested in the definition of methods to translate various kinds of data sources into RDF. R2RML [2], for instance, is the W3C recommendation to describe RDB-to-RDF mappings. RML extends R2RML for the integration of heterogeneous data formats [3]. To our knowledge though, no method has been proposed yet to tackle NoSQL-to-RDF translation.

In this paper, we present xR2RML, a mapping language designed as an extension of R2RML and RML. Besides relational databases, xR2RML addresses the mapping of a large and extensible scope of non-relational databases to RDF. It is designed to flexibly adapt to various data models and query languages, it can translate data with mixed formats and generate RDF collections and containers. xR2RML is exemplified and validated through a real-world use case. A large-scale SKOS thesaurus aimed at supporting studies on History of Zoology, Archaeozoology and Conservation Biology is thus generated.

In the rest of this section we draw a picture of other works pertaining to the translation of various data sources to RDF, and we scope the objectives of xR2RML. Section 2 explores in more details the capabilities required to reach these goals. In section 3 we recall the main characteristics of R2RML and RML, and in section 4 we describe xR2RML specific extensions. Section 5 presents a working implementation of the language. Section 6 describes the experimentation we ran to create a large taxonomical reference in SKOS. Finally sections 7 and 8 discuss xR2RML applicability in different contexts and concludes by outlining some perspectives.

1.1 Related Works

Wrapper-based data integration systems like Garlic [4] and SQL/MED [5] generally have similar architectures: a global data model is described using specific modelling languages (e.g. Garlic’s GDL), a query federation engine handles user queries expressed in terms of a global data model and determines a query plan, a per-data source wrapper implements a specific wrapper interface and performs the mapping with the data source schema. No guideline is provided as to how a wrapper should describe and implement the mapping.

The same global architecture holds in data integration systems based on semantic web technologies. Existing works focus on efficient query planning and distribution, such as FedX [6], Anapsid [7] and KGRAM-DQP [8]. The global data model is expressed by domain ontologies using common languages, e.g.

RDFS or OWL. User queries, expressed in terms of the domain ontologies, are written in SPARQL. SPARQL is also used as the wrapper interface. Each data source wrapper is a SPARQL endpoint that performs the schema mapping with the source schema. Our work, as well as most related works listed below, focuses on the mapping step: the rationale is to standardize the schema mapping description, so that a mapping description can be written once and applied with different wrapper implementations.

RDB-to-RDF mapping has been an active field of research during the last ten years [9,10,11]. Several mapping methods and languages have been proposed over time, based either on the materialization of RDF data sets or on the SPARQL-based access to relational data. Published in 2012, R2RML, the W3C RDB-to-RDF mapping language recommendation, has reached a notable consensus³.

Similarly, various solutions exist to map XML data to RDF. The XSPARQL query language [12] combines XQuery and SPARQL for bidirectional transformations between XML and RDF. Several other solutions are based on the XSLT technology such as XML Scissor-lift [13] that describes mapping rules in Schema-tron XML validation language, and AstroGrid-D [14]. SPARQL2XQuery [15] applies XML Schema to RDF/OWL translation rules.

Much work has already been accomplished regarding the translation of CSV, TSV and spreadsheets to RDF. Tools have been developed such as XLWrap [16] and RDF Refine⁴. The Linked CSV⁵ format is a proposition to embed metadata in a CSV file, that makes it easy to link on the Web and eventually to translate to RDF or JSON. However this approach assumes that CSV data is made compliant with the format in the first place, before it can be translated to RDF. The CSV on the Web W3C Working Group⁶, created in 2014, intends to propose a recommendation for the description of and access to CSV data on the Web. In this context, RDF is one of the formats targeted either to represent metadata about CSV data, or as a format to translate CSV data into.

Several tools are designed as frameworks for the integration of sources with heterogeneous data formats. XSPARQL, cited above, provides an R2RML-compliant extension. Thus it can simultaneously translate relational, XML and RDF data to XML or RDF. TARQL⁷ is a SPARQL-based mapping language that can convert from RDF, CSV/TSV and JSON formats to RDF, but it does not focus on how the data is retrieved from different types of databases. Datalift [17] provides an integrated set of tools for the publication in RDF of raw structured data (RDB, CSV, XML) and the interlinking of resulting data sets.

RML [18,3] is an extension of R2RML that tackles the mapping of data sources with heterogeneous data formats such as CSV/TSV, XML or JSON. Most approaches create links between data sets after they were translated to RDF, e.g. using properties `rdfs:seeAlso` or `owl:sameAs`. This is sometimes not

³ <http://www.w3.org/2001/sw/rdb2rdf/wiki/Implementations>

⁴ <http://refine.deri.ie/>

⁵ <http://jenit.github.io/linked-csv/>

⁶ <http://www.w3.org/2013/csvw/wiki>

⁷ <https://github.com/cygri/tarql/wiki/TARQL-Mapping-Language>

adequate as logical resources having different identifiers in different data sets cannot easily be reconciled. RML creates linked data sets at mapping time by enabling the simultaneous mapping of multiple data sources, thus allowing for cross-references between resources defined in various data sources. However, RML does not investigate the constraints that arise when dealing with different types of databases. It proposes a solution to reference data elements within query results using expressive languages such as XPath and JSONPath. But it does not clearly distinguish between such languages and the actual query language of a database. In some cases they might be the same, e.g. XPath can be used to query an XML native database, and later on to reference data elements from query results. But in the general case, the query language and the language used to reference elements within query results must be dissociated, e.g. NoSQL document stores use proprietary query languages, while results are JSON documents that can be evaluated against JSONPath expressions. Furthermore, RML explicitly refers to known evaluation languages (ql:JSONPath, ql:XPath). In this context, supporting a new evaluation language requires to change the mapping language definition. To achieve more flexibility, we believe that such characteristics should be implementation-dependent, leaving the mapping language free from any explicit dependency.

1.2 Objectives of this Work

The works presented in section 1.1 address various types of data sources. Some of them could be extended to new data sources by developing ad-hoc extensions, although they are generally not designed to easily support new data models and query languages. Only RML comes with this flexibility as its design aims at adapting to new data models. Our goal with xR2RML is to define a generic mapping language able to equally apply to most common relational and non-relational databases. We make a specific focus on NoSQL and XML native databases, and we argue that our work can be generalized to some other types of database, for instance object-oriented and directory (LDAP) databases. In section 2 we explore the capabilities required by xR2RML to reach these goals.

Moreover, we describe a validation of xR2RML made in the context of a real-world use case: the translation of data from a MongoDB NoSQL document store into a large-scale RDF-based thesaurus aimed at supporting studies on History of Zoology, Archaeozoology and Conservation Biology.

2 xR2RML language requirements

Different kinds of databases typically differ in several aspects: the query language used to retrieve data, the data model that underlies the data structures retrieved and the cross-data referencing scheme, if any. Below we explore in further details the capabilities that we want xR2RML to provide.

Query languages. The landscape of modern database systems shows a vast diversity of query languages. Relational databases generally support ANSI

SQL, and most native XML databases support XPath and XQuery. By contrast, NoSQL is a catch-all term referring to very diverse systems [19,20]. They have heterogeneous access methods ranging from low-level APIs to expressive query languages. Despite several propositions of common query language (N1QL⁸, UnQL⁹, SQL++ [21], ArangoDB QL¹⁰, CloudMdsQL [22]), no consensus has emerged yet, that would fit most NoSQL databases. Therefore, until a standard eventually arises, xR2RML must be agile enough to cope with various query languages and protocols in a transparent manner.

Data models. Similarly to the case of query languages, we observe a large heterogeneity in data models of modern databases. To describe their translation to RDF, a mapping language must be able to reference any data element from their data models. Below we list most common data models, we shortly analyse formats in which data is retrieved and figure out how a mapping language can reference data elements within retrieved data.

Relational databases comply with a row-based model in which column names uniquely reference cells in a row. NoSQL extensible column stores¹¹ also comply with the row-based model, with the difference that all rows do not necessarily share the same columns. For such systems, referencing data elements is simply achieved using column names. Other non-relational systems, such as XML native databases, NoSQL key-value stores, document stores or graph stores, have heterogeneous data models that can hardly be reduced to a row-based model:

- In databases relying on a specific data representation format like JSON (notably in NoSQL document stores) and XML, data is stored and retrieved as documents consisting of tree-like compound values. Referencing data elements within such documents can be achieved thanks to languages such as JSONPath and XPath.
- Object-oriented databases conventionally provide methods to serialize objects, typically as key-value associations: keys are attribute names while values are objects (composition or aggregation relationship), or compound values (collection, map, etc). Serialization is typically done in XML or JSON, thus here again we can apply XPath or JSONPath expressions.
- A directory data model is organised as a tree: each node has an identifier and a set of attributes represented as **name=value**. Each entry retrieved from an LDAP request is named using an LDAP path expression, e.g. **cn=Franck Michel,ou=cnrs,o=fr**. Referencing data elements within such entries can be simply achieved using attribute names.
- In graph databases, the abstract data model basically consists of nodes and edges. Query capabilities generally allow to retrieve either values matching specific patterns (like the SPARQL SELECT clause), or a set of nodes and edges representing a result graph (like the SPARQL CONSTRUCT clause). Whatever the type of result though, graph databases commonly provide

⁸ <http://www.couchbase.com/communities/n1ql>

⁹ <http://unql.sqlite.org/index.html>

¹⁰ <http://docs.arangodb.org/Aql/README.html>

¹¹ aka. column family store, column-oriented store, etc.

APIs to manipulate query results. For instance a SPARQL SELECT result set has a row-based format: each row of a result set consists of columns typically named after query variable names. The Neo4J graph database provides a JDBC interface to process a query result, and its REST interface returns result graphs as JSON documents. Thus, although a graph may be a somehow complex data structure, query results can be fairly easy to manipulate using well-known formats: a row-column model, a serialization in JSON or some other representation syntax, etc.

Finally, the way a mapping language can reference data elements within query results depends more on the API capabilities than the data model itself. To be effective, xR2RML must transparently accept any type of data element reference expression. This includes a column name (applicable not only to row-based data models but also to any row-based query result), JSONPath, XPath or LDAP path expressions, etc. An xR2RML processing engine must be able to evaluate such expressions against query results, but the mapping language itself must remain free from any reference to specific expression syntaxes.

Collections. Many data models support the representation of collections: these can be sets, arrays or maps of all kinds (sorted or not, with or without duplicates, etc.). Although the RDF data model supports such data structures, to the best of our knowledge, existing mapping languages do not allow for the production of RDF collections (`rdf:List`) nor RDF containers (`rdf:Bag`, `rdf:Seq`, `rdf:Alt`), except TARQL that is able to convert a JSON array into an `rdf:List`. In all other cases, structured values such as collections or key-value associations are flattened into multiple RDF triples. Listing 1.1 is an example XML collection consisting of two “movie” elements.

Its translation into two triples is illustrated in Listing 1.2. Assuming that the order of “movie” elements implicitly represents the chronological order in which movies were shot, triples in Listing 1.2 lose this information. Using an RDF sequence may be more appropriate in this case, as illustrated in Listing 1.3.

```
<director name="Woody Allen">
  <movie>Annie Hall</movie>
  <movie>Manhattan</movie>
</director>
```

Listing 1.1. Example of XML collection

```
<http://example.org/dir/Woody\%20Allen>
  ex:directed "Annie Hall".
<http://example.org/dir/Woody\%20Allen>
  ex:directed "Manhattan".
```

Listing 1.2. Translation to multiple RDF triples

```
<http://example.org/dir/Woody\%20Allen>
ex:movieList [ a rdf:Seq;
  rdf:_1 "Annie Hall"; rdf:_2 "Manhattan" ].
```

Listing 1.3. Translation to an RDF sequence

Consequently, to map heterogeneous data to RDF while preserving concepts such as collections, bags, alternates or sequences, xR2RML must be able to map data elements to RDF collections and containers.

Cross-references. Cross-references are commonly implemented as foreign key constraints in relational data models, or aggregation and composition relationships in object-oriented models. Cross-referencing is even the primary goal of graph-based databases. More generally, it is possible to cross-reference logical entities in any type of database. For instance, a JSON document of a NoSQL document store may refer to another document by its identifier or any other field that identifies it uniquely, even if this is generally not recommended for the sake of performances.

A cross-referenced logical resource may be mapped alternatively as the subject or the object of triples. This may entail joint queries between tables or documents. Therefore, xR2RML must (i) allow a modular description so that the mapping of a logical resource can be written once and easily reused as a subject or an object, and (ii) allow the description of joint queries to retrieve cross-referenced logical resources.

Summary. Finally, we draw up the list of key capabilities expected from xR2RML as follows:

1. It enables to describe the mapping of various relational and non-relational databases to RDF.
2. It is flexible enough to allow for new databases, query languages and data models in an agile manner: supporting a new system, query language and/or data model only requires changes in the implementation (adaptor, plug-in, etc.), but no changes are required in the mapping language itself.
3. It enables to generate RDF collections (`rdf:List`) or containers (`rdf:Seq`, `rdf:Bag`, `rdf:Alt`) from one-to-many relations modelled as compound values or as cross-references. RDF collections and containers can be nested.
4. It enables to perform joint queries following cross-references between logical resources, and it allows the modular reuse of mapping definitions.

Additionally, data sources to be mapped to RDF using xR2RML should provide a **declarative** query language. If not, it must be possible to fetch the whole data at once, like a CSV or XML file returned by a Web service. There must exist technical means to parse query results, ranging from simple column names to expressive languages like XPath. In case of large data sets, the database interface should provide ways to iterate on query results, similarly to SQL cursors in RDBs.

To help in the design of xR2RML we chose to leverage R2RML, a standard, well-adopted mapping language for relational databases. R2RML already provides some of the requirements listed above: modularity, management of cross-references, as well as rich features such as the ability to define target named graphs. To facilitate its understanding and adoption, xR2RML is designed as a backward compatible extension of R2RML. Besides, to address the mapping of heterogeneous data formats such as CSV/TSV, XML and JSON, we leverage propositions of RML that is itself an extension of R2RML.

3 R2RML and RML

R2RML is a generic language meant to describe customized mappings that translate data from a relational database into an RDF data set. An R2RML mapping is expressed as an RDF graph written in Turtle syntax¹². An R2RML mapping graph consists of *triples maps*, each one specifying how to map rows of a logical table to RDF triples. A triples map is composed of exactly one *logical table* (property `rr:logicalTable`), one *subject map* (property `rr:subjectMap`) and any number of *predicate-object maps* (property `rr:predicateObjectMap`). A logical table may be a table, an SQL view (property `rr:tableName`), or the result of a valid SQL query (property `rr:sqlQuery`). A predicate-object map consists of *predicate maps* (property `rr:predicateMap`) and *object maps* (property `rr:objectMap`). For each row of the logical table, the subject map generates a subject IRI, while each predicate-object map creates one or more predicate-object pairs. Triples are produced by combining the subject IRI with each predicate-object pair. Additionally, triples are generated either in the default graph or in a named graph specified using *graph maps* (property `rr:graphMap`).

Subject, predicate, object and graph maps are all R2RML *term maps*. A term map is a function that generates RDF terms (either a literal, an IRI or a blank node) from elements of a logical table row. A term map must be exactly one of the following: a *constant-valued term map* (property `rr:constant`) always generates the same value; a *column-valued term map* (property `rr:column`) produces the value of a given column in the current row; a *template-valued term map* (property `rr:template`) builds a value from a template string that references columns of the current row.

When a logical resource is cross-referenced, typically by means of a foreign key relationship, it may be used as the subject of some triples and the object of some others. In such cases, a *referencing object map* uses IRIs produced by the subject map of a (parent) triples map as the objects of triples produced by another (child) triples map. In case both triples maps do not share the same logical table, a joint query must be performed. A join condition (property `rr:joinCondition`) names the columns from the parent and child triples maps, that must be joined (properties `rr:parent` and `rr:child`).

Below we provide a short illustrative example. Triples map `<#R2RML_Directors>` uses table `DIRECTORS` to create triples linking movie directors (whose IRIs are built from column `NAME`) with their birth date (column `BIRTH_DATE`).

```
<#R2RML_Directors>
rr:logicalTable [ rr:tableName "DIRECTORS" ];
rr:subjectMap [
  rr:template "http://example.org/dir/{NAME}" ];
rr:predicateObjectMap [
  rr:predicate ex:birthdate;
  rr:objectMap [
    rr:column "BIRTH_DATE"; rr:datatype xsd:date ] ].
```

¹² <http://www.w3.org/TR/turtle/>

RML is an extension of R2RML that targets the simultaneous mapping of heterogeneous data sources with various data formats, in particular hierarchical data formats. An RML logical source (property `rml:logicalSource`) extends R2RML logical table and points to the data source (property `rml:source`): this may be a file on the local file system, or data returned from a web service for instance. A reference formulation (property `rml:referenceFormulation`) names the syntax used to reference data elements within the logical source. As of today, possible values are `q1:JSONPath` for JSON data, `q1:XPath` for XML data, and `rr:SQL2008` for relational databases. Data elements are referenced with property `rml:reference` that extends `rr:column`. Its object is an expression whose syntax matches the reference formulation. Similarly, the definition of property `rr:template` is extended to allow such reference expressions to be enclosed within curly braces (`'{'` and `'}'`). Below we provide an RML example. It is very similar to the R2RML example above, with the difference that data now comes from a JSON file “directors.json”.

```
<#RML_Directors>
rml:logicalSource [
  rml:source "directors.json";
  rml:referenceFormulation q1:JSONPath;
  rml:iterator "$.*"; ]
rr:subjectMap [
  rr:template "http://example.org/dir/{$.*.name}" ];
rr:predicateObjectMap [
  rr:predicate ex:birthdate;
  rr:objectMap [
    rml:reference "$.*.birthdate"; rr:datatype xsd:date ] ].
```

4 The xR2RML Mapping Language

In this section we briefly describe the elements of the xR2RML language. A complete specification is provided in [23]. We illustrate the descriptions with a running example: Listing 1.4 shows JSON documents stored in a MongoDB database, in two collections: a “directors” collection with documents on movie directors, and a “movies” collection in which movies are grouped in per-decade documents. Listing 1.5 shows an xR2RML mapping graph to translate those documents into RDF. Director IRIs are built using director names, while movie IRIs use movie codes. We assume the following namespace prefix definitions:

```
@prefix xrr: <http://www.i3s.unice.fr/ns/xr2rml#>.
@prefix rr: <http://www.w3.org/ns/r2rml#>.
@prefix rml: <http://semweb.mmlab.be/ns/rml#>.
@prefix xsd: <http://www.w3.org/2001/XMLSchema#>.
@prefix ex: <http://example.com/ns#>.
```

4.1 Describing A Logical Source

To reach its genericity objective, xR2RML must avoid explicitly referring to specific query languages or data models. Keeping this in mind, we define logical sources as a mean to represent a data set from any kind of database. In conformance with R2RML principles, we keep database connection details out of the scope of the mapping language. In RML on the other hand, a logical source points to the data to be mapped typically using a file URL (property `rml:source`). This difference makes it difficult for xR2RML to extend RML's logical source concept. Instead, xR2RML extends the R2RML logical source while commonalities are addressed by using or extending some RML properties (`rml:referenceFormulation`, `rml:query`, `rml:iterator`).

xR2RML triples maps extend R2RML triples maps by referencing a *logical source* (property `xrr:logicalSource`) which is the result of a request applied to the input database. It is either an *xR2RML base table* or an *xR2RML view*. The xR2RML base table extends the concept of *R2RML table or view* to tabular databases beyond relational databases (extensible column store, CSV/TSV, etc.). It refers to a table by its name (property `rr:tableName`). An xR2RML view represents the result of executing a query against the input database. It has exactly one `xrr:query` property that extends RML property `rml:query` (which itself extends `rr:sqlQuery`¹³). Its value is a valid expression with regards to the query language supported by the input database. No assumption is made whatsoever as to the query language used.

Reference formulation. Retrieving values from a query result set requires evaluating *data element references* against the query result. Relational database APIs (such as JDBC drivers) support the evaluation of a column name against the current row of a result set. Conversely, some databases come with simple APIs that provide lower level evaluation features. For instance, APIs of most NoSQL document stores return JSON documents but hardly support JSON-Path. Therefore, the xR2RML processing engine is responsible for evaluating such data element references. To do so, it needs to know which syntax is being used. To this end, RML introduced the reference formulation concept (property `rml:referenceFormulation` of a logical source) to name the syntax of data element references. As underlined above, xR2RML adheres to R2RML's principle that database-specific details be kept out of the scope of the mapping language. We also want the mapping language to remain free from explicit reference to specific syntaxes. As a result, we amend the R2RML processor definition as follows: *an xR2RML processor must be provided with a database connection and the reference formulation applicable to results of queries run against the connection. If the reference formulation is not provided, it defaults to column name, in order to ensure backward compatibility with R2RML.*

Iteration model. In R2RML, the row-based iteration occurs on a set of rows read from a logical table. xR2RML applies this principle to other systems

¹³ `rml:query` also subsumes `rml:xmlQuery` and `rml:queryLanguage`, although none of those properties are described or exemplified in the RML language specification and articles at the time of writing.

```

Collection "directors":
{"name": "Woody Allen", "directed":
  ["Manhattan", "Interiors"]},
{"name": "Wong Kar-wai", "directed":
  ["2046", "In the Mood for Love"]}

Collection "movies":
{ "decade": "2000s", "movies": [
  {"name": "2046", "code": "m2046",
    "actors": ["T. Leung", "G. Li"]},
  {"name": "In the Mood for Love", "code": "Mood",
    "actors": ["M. Cheung"]} ] }
{ "decade": "1970s", "movies": [
  {"name": "Manhattan", "code": "Manh",
    "actors": ["Woody Allen", "Diane Keaton"]}
  {"name": "Interiors", "code": "Int01",
    "actors": ["D. Keaton", "G. Page"]} ] }

```

Listing 1.4. Example Database

```

<#Movies>
  xrr:logicalSource [
    xrr:query "db.movies.find({decade:{$exists:true}})";
    rml:iterator "$.movies.*";
  ];
  rr:subjectMap [
    rr:template "http://example.org/movie/{$.code}" ];
  rr:predicateObjectMap [
    rr:predicate ex:starring;
    rr:objectMap [
      rr:termType xrr:RdfBag;
      xrr:reference "$.actors.*";
      xrr:nestedTermMap [ rr:datatype xsd:string ] ] ].

<#Directors>
  xrr:logicalSource [ xrr:query "db.directors.find()" ];
  rr:subjectMap [
    rr:template "http://example.org/dir/{$.name}" ];
  rr:predicateObjectMap [
    rr:predicate ex:directed;
    rr:objectMap [
      rr:parentTriplesMap <#Movies>;
      rr:joinCondition [
        rr:child "$.directed.*";
        rr:parent "$.name" ] ] ].

```

Listing 1.5. xR2RML Example Mapping Graph

returning row-based result sets: CSV/TSV files, extensible column stores, but also some graph databases as underlined in 2, e.g. a SPARQL SELECT result set is a table in which columns are named after the variables in the SELECT clause. In the context of non row-based result sets, the model is implicitly extended to a *document-based iteration model*: a document is basically one entry of a result set returned by the database, e.g. a JSON document retrieved from a NoSQL document store, or an XML document retrieved from an XML native database. In the case of data sources whose access interface does not provide built-in iterators, e.g. a web service returning an XML response at once, then a single iteration occurs on the whole retrieved document.

Yet, some specific needs may not be fulfilled. For instance, it may be needed to iterate on explicitly specified entries of a JSON document or elements of an XML tree. To this end, we leverage the concept of iterator introduced in RML. An iterator (property `rml:iterator`) specifies the iteration pattern to apply to data read from the input database. Its value is a valid expression written using the syntax specified in the reference formulation. The iterator can be either omitted or empty when the reference formulation is a column name.

Listing 1.5 presents two logical source definition examples. Both consist of a MongoDB query (property `xrr:query`). We assume that the JSONPath reference formulation is provided along with the database connection. In collection “directors” (Listing 1.4), each document describes exactly one director. By contrast, in collection “movies” each document refers to several movies grouped by decade. To avoid mixing up multiple movies of a single document, an iterator with JSONPath expression `$.movies.*` is associated with triples map `<#Movies>`: thus, the triples map applies separately on each movie of each document.

4.2 Referencing Data Elements

In section 3 we have seen that RML properties `rml:reference` and `rr:template` both allow data element references expressed according to the reference formulation (column name, XPath, JSONPath). xR2RML uses these RML definitions as a starting point to a broader set of use cases.

In real world use cases, databases commonly store values written in a data format that they cannot interpret. For instance, in key-value stores and in most extensible column stores, values are stored as binary objects whose content is opaque to the system. A developer may choose to embed JSON, CSV or XML values in the column of a relational table, for performance issues or due to application design constraints. We call such cases *mixed content*.

xR2RML proposes to apply the principle of data element references defined in RML, and extend it to allow referencing data elements within mixed content. An xR2RML *mixed-syntax path* consists of the concatenation of several path expressions, each path being enclosed in a *syntax path constructor* that explicits the path syntax. Existing constructors are: `Column()`, `CSV()`, `TSV()`, `JSONPath()` and `XPath()`. For example, in a relational table, a text column `NAME` stores JSON-formatted values containing people’s first and last names, e.g.: `{"First": "John",`

"Last":"Smith"}). Field `FirstName` can be referenced with the following mixed-syntax path: `Column(NAME)/JSONPath($.First)`. An xR2RML processing engine evaluates a mixed-syntax path from left to right, passing the result of each path constructor on to the next one. In this example, the first path retrieves the value associated with column `NAME`. Then the value is passed on to the next path constructor that evaluates `JSONPath` expression `"$.First"` against the value. The resulting value is finally translated into an RDF term according to the current term map definition.

xR2RML defines property `xrr:reference` as an extension of RML property `rml:reference`, and extends the definition of property `rr:template`. Both properties accept either simple references (illustrated in Listing 1.5) or mixed-syntax path expressions.

4.3 Producing RDF Terms and (Nested) RDF Collections/Containers

In a row-based logical source, a valid column name reference returns zero or one value during each triples map iteration. In turn an R2RML term map generates zero or one RDF term per iteration. By contrast, `JSONPath` and `XPath` expressions used with properties `xrr:reference` and `rr:template` allow addressing multiple values. For instance, `XPath` expression `//movie/name` returns all `<name>` elements of all `<movie>` elements. Therefore, reference-valued and template-valued term maps can return multiple RDF terms at once. This difference entails the definition of two strategies with regards to how triples maps combine RDF terms to build triples: the product strategy, and the collection/container strategy.

Product strategy. During each iteration of an xR2RML triples map, triples are generated as the product between RDF terms produced by the subject map and each predicate-object pair. Predicate-object pairs result of the product between RDF terms produced by the predicate maps and object maps of each predicate-object map. Like any other term map, a graph map may also produce multiple terms. The product strategy equally applies in that case, therefore triples are produced simultaneously in all target graphs corresponding to the multiple RDF terms produced by the graph map.

Collection/container strategy. Multiple values returned by properties `xrr:reference` and `rr:template` are combined into an RDF collection or container. This is achieved using new xR2RML values of the `rr:termType` property: a term map with term type `xrr:RdfList` generates an RDF term of type `rdf:List`, term type `xrr:RdfSeq` corresponds to `rdf:Seq`, `xrr:RdfBag` to `rdf:Bag` and `xrr:RdfAlt` to `rdf:Alt`. Listing 1.5 illustrates this case: instead of generating multiple triples relating each movie to one actor, triples map `<#Movies>` relates each movie to a bag of actors starring in that movie. For instance:

```
<http://example.org/movie/m2046> ex:starring
  [ a rdf:Bag; rdf:_1 "Tony Leung"; rdf:_2 "Gong Li" ].
```

At this point, two important needs must still be addressed in the collection/container strategy: (i) like in a regular term map, it must be possible to assign a term type, language tag or data type to the members of an RDF collection or

container; and (ii) it must be possible to nest any number of RDF collections and containers inside each-other. Both needs are fulfilled using *xR2RML Nested Term Maps*. A nested term map (property `xrr:nestedTermMap`) very much resembles a regular term map, with the exception that it can be defined only in the context of a term map that produces RDF collections or containers. In a column-valued or reference-valued term map, a nested term map describes how to translate values read from the logical source into RDF terms, by specifying optional properties `rr:termType`, `rr:language` and `rr:datatype`. Similarly, in a template-valued term map, a nested term map applies to values produced by applying the template string to input values. In Listing 1.5, triples map `<#Movie>` uses a nested term map to assign a `xsd:string` datatype to names of names starring in a movie. For instance:

```
<http://example.org/movie/m2046> ex:starring [ a rdf:Bag;
  rdf:_1 "Tony Leung"^^xsd:string; rdf:_2 "Gong Li"^^xsd:string ] .
```

Finally, properties `xrr:reference` and `rr:template` can be used within a nested term map to recursively parse structured values while producing nested RDF collections and containers.

4.4 Reference Relationships Between Logical Sources

A cross-referenced logical resource usually serves as the subject of some triples and the object of other triples. In R2RML, this is achieved using a referencing object map. xR2RML extends R2RML referencing object maps in two ways. Firstly, when a joint query is needed (i.e. the parent and child triples map do not share the same logical source), properties `rr:child` and `rr:parent` of the join condition contain data element references (4.2), possibly including mixed-syntax paths. As underlined in section 4.3, such data element references may produce multiple terms. Consequently, the equivalent joint query of a referencing object map must deal with multi-valued child and parent references. More precisely, a join condition between two multi-valued references should be satisfied if at least one data element of the child reference matches one data element of the parent reference. This is described in Definition 1 using an SQL-like syntax and first order logic for the description of WHERE conditions.

Definition 1: If a referencing object map has at least one join condition, then its equivalent joint query is:

$$\begin{aligned}
 & \text{SELECT * FROM (child-query) AS child, (parent-query) AS parent} \\
 & \quad \text{WHERE} \\
 & \quad \exists c_1 \in \text{eval}(\text{child}, \{\text{child-ref}_1\}), \exists p_1 \in \text{eval}(\text{parent}, \{\text{parent-ref}_1\}), c_1 = p_1 \\
 & \quad \text{AND} \\
 & \quad \exists c_2 \in \text{eval}(\text{child}, \{\text{child-ref}_2\}), \exists p_2 \in \text{eval}(\text{parent}, \{\text{parent-ref}_2\}), c_2 = p_2 \\
 & \quad \text{AND ...}
 \end{aligned}$$

where “{child-ref *i*}” and “{parent-ref *i*}” are the child and parent references of the *i*th join condition, and “eval(child, {ref})” and “eval(parent, {ref})” are the result of evaluating data element reference “{ref}” on the result of the child and parent queries.

Listing 1.5 depicts a simple example: in triples map `<#Directors>`, the object map uses movie IRIs generated by parent triples map `<#Movies>`. When processing director “Wong Kar-wai”, the child reference (`$.directed.*`) returns values “2046” and “In the Mood for Love”, while the parent reference (`$.name`) returns a single movie name. The join condition is satisfied if the parent reference returns one of “2046” or “In the Mood for Love”. Generated triples use movie codes to build movie IRIs, such as:

```
<http://example.org/dir/Wong%20Kar-wai>
  ex:directed <http://example.org/movie/m2046>.
```

Secondly, the objects produced by a referencing object map can be grouped in an RDF collection or container, instead of being the objects of multiple triples. To do so, an xR2RML referencing object map may have a `rr:termType` property with value `xrr:RdfList`, `xrr:RdfSeq`, `xrr:RdfBag` or `xrr:RdfAlt`. Results of the joint query are grouped by child value, i.e. objects generated by the parent triples map, referring to the same child value, are grouped as members of an RDF collection or container. An interesting consequence of this use case is the ability, in the case of a regular relational database, to build an RDF collection or container reflecting a one-to-many relation.

5 Implementation

To evaluate the effectiveness of xR2RML, we have developed Morph-xR2RML, an open source prototype implementation available on Github¹⁴. It is written in Scala and is an extension of Morph-RDB [24], an R2RML implementation.

In a first step, we upgraded Morph-RDB to support xR2RML features in the context of relational databases. This included the support of logical sources, mixed contents (JSON, XML, CSV or TSV data embedded in cells) and RDF collections/containers. In a second step, we developed a connector to the MongoDB document store, to translate MongoDB JSON documents into RDF. A MongoDB query string is specified in each triples map logical source. The connector executes the query and iterates over result documents returned by the database. Subsequently, results are passed to the xR2RML processor that applies the optional iterator (`rml:iterator`) and evaluates JSONPath expressions in each `xrr:reference` and `rr:template` property of all term maps. The support of RDF collections/containers was validated, including in the case of referencing object maps that entail a joint query between two JSON documents.

We evaluated the prototype using two simple databases: a MySQL relational database and a MongoDB database with two collections. In both cases, the data and associated xR2RML mappings were written to cover most mapping situations addressed by xR2RML: strategies for handling multiple RDF terms, mixed-syntax paths with mixed contents (relational, JSON, XML, CSV/TSV), cross-references, RDF collection/containers, UTF-8 encoding. Both databases as well as the example mappings are available on the GitHub repository. The current status of the prototype applies the data materialization approach, i.e. RDF

¹⁴ <https://github.com/frmichel/morph-xr2rml/>

data is generated by sequentially applying all triples maps. The query rewriting approach (SPARQL to database specific query rewriting) may be considered in future work as suggested in section 8. At the time of writing the prototype has two limitations: (i) only one level of RDF collections and containers can be generated (no nested collections/containers), and (ii) the result of a joint query in a relational database cannot be translated into an RDF collection or container.

6 Validation: construction of a SKOS zoological and botanical reference thesaurus

The Zoomathia research network¹⁵ studies the transmission of zoological knowledge throughout historical periods. It intends to leverage the Semantic Web technologies to annotate and link together various resources such as rich medieval compilation literature on Ancient zoological knowledge, archaeozoological data from excavation reports, iconographic material and modern conservation biology knowledge. This challenging goal can be addressed through the use of controlled and widely accepted semantic references. In this context, the TAXREF [25] zoological and botanical taxonomy has been chosen to build a SKOS thesaurus¹⁶ supporting the integration of these heterogeneous data sets. SKOS, the Simple Knowledge Organization System, is a W3C standard designed to represent controlled vocabularies, taxonomies and thesauri. It is extensively used to bridge the gap between existing knowledge organisation systems and the Semantic Web and Linked Data. In this section we shortly present TAXREF and we describe how we used xR2RML for the creation of a SKOS vocabulary faithfully representing TAXREF. More detailed information about TAXREF's content and structure, the SKOS modeling and the data sources to be integrated, can be found in [26].

TAXREF is the French national taxonomic reference for fauna, flora and fungus of metropolitan France and overseas departments and collectivities. It registers 452.106 taxa of living beings from the Palaeolithic until now, covering continental and marine environments. Each taxon is provided along with its scientific name, upper taxon in the classification, synonyms, vernacular names, authority (author name and publication year), taxonomical rank (order, family, gender, species...), type of habitat (marine, terrestrial...) and a biogeographical status (present, endemic, extinct, etc.) for each considered geographical area. TAXREF is developed, maintained and distributed by the French *National Museum of Natural History* (MNHN). It can be browsed, downloaded in TSV format, or queried through a Web service.

In a first step we defined the SKOS modelling of TAXREF¹⁷. In brief, a SKOS concept is created for each taxon, along with a SKOS label for each scientific (preferred label) name and synonym (alternate label). The SKOS broader property is used to model the relationships between a taxon and the upper taxon in

¹⁵ <http://www.cepam.cnrs.fr/zoomathia/>

¹⁶ <http://www.w3.org/2009/08/skos-reference/skos.html>

¹⁷ Changes were made since the description presented in [26]

the classification. To ensure proper linkage with well-adopted data sets, in particular within the Linking Open Data cloud, we identified relevant ontologies such as the NCBI taxonomic classification¹⁸, the GeoSpecies ontology¹⁹, the ENVO²⁰ environment ontology and Biodiversity Information Standards promoted by the Taxonomic Databases Working Group²¹. Links were made either by using appropriate properties and classes, or by aligning taxa with their equivalent in other ontologies.

Listing 1.6 shortly illustrates the SKOS representation of a dolphin species, the taxon "Delphinus delphis", generated by xR2RML using the Turtle RDF syntax. This consists of a SKOS concept for the taxon and two SKOS labels, one for the reference name and one for the synonym "Delphinus tropicalis".

In a second step we retrieved a full TAXREF JSON dump from the TAXREF Web service, and imported it into a MongoDB instance. We wrote the xR2RML mappings that describe how to map the result of queries to the MongoDB instance into RDF triples, according to the SKOS modelling. This step requires a good knowledge of the xR2RML language, but it is quite straightforward for users who are already familiar with R2RML.

To perform the translation of TAXREF into SKOS, we ran Morph-xR2RML, the prototype implementation of xR2RML described in section 5, on a laptop equipped with a 3GHz Intel Core i7 processor and 8GB of RAM. The resulting RDF graph consists of more than 5 million triples. The translation process required approximately 6 hours to complete, which can be considered surprisingly long. We analyze this issue in the next paragraph. Nonetheless, this execution time span should not be considered as an hurdle in the context of TAXREF. Indeed, insofar as TAXREF is updated once a year approximately, the traditional Extract, Transform and Load (ETL) approach is relevant. This is the approach we intend to follow in the future: each time TAXREF is updated, the corresponding SKOS thesaurus is generated. The resulting graph is loaded into a public triple store accessible using either a SPARQL endpoint or the HTTP dereferencing method.

The xR2RML mapping graph for TAXREF consists of 90 triples maps. The high number of triples maps is a consequence of the distance between the internal structure of TAXREF and the targeted SKOS modelling. We illustrate this distance with an example. Habitats are coded in TAXREF with integer values, e.g. value '1' represents the marine habitat. Translating the marine habitat into URI `<http://inpn.mnhn.fr/taxref/habitat#1>` would be straightforward: a template-valued term map could simply append the value '1' read from the database to the namespace `<http://inpn.mnhn.fr/taxref/habitat#>`. Therefore a single triples map (thus a single query) would be sufficient to generate all triples related to all types of habitat. However, we wish to generate more meaningful URIs, such as `<http://inpn.mnhn.fr/taxref/habitat#Marine>`. This URI cannot

¹⁸ <http://www.ontobee.org/browser/index.php?o=NCBITaxon>

¹⁹ <http://datahub.io/dataset/geospecies>

²⁰ <http://www.ontobee.org/browser/index.php?o=ENVO>

²¹ <http://www.tdwg.org/>

```

@prefix rdfs: <http://www.w3.org/2000/01/rdf-schema#> .
@prefix dct: <http://purl.org/dc/elements/1.1/> .
@prefix skc: <http://www.w3.org/2004/02/skos/core#>.
@prefix skx: <http://www.w3.org/2008/05/skos-xl#>.
@prefix tc: <http://lod.taxonconcept.org/ontology/txn.owl#>.
@prefix nt: <http://purl.obolibrary.org/obo/ncbitaxon#> .
@prefix dwc: <http://rs.tdwg.org/dwc/terms/> .
@prefix taxr: <http://inpn.mnhn.fr/taxref/> .

<http://inpn.mnhn.fr/taxref/8.0/taxon/60878> a skc:Concept;
skx:altLabel <http://inpn.mnhn.fr/espece/cd_nom/60881>;
skx:prefLabel <http://inpn.mnhn.fr/espece/cd_nom/60878>;
skc:broadier <http://inpn.mnhn.fr/taxref/taxon/191591>;
taxr:hasHabitat <http://inpn.mnhn.fr/taxref/habitat#Marine>;
nt:has_rank <http://inpn.mnhn.fr/taxref/taxrank#Species>;
skc:note "Delphinus delphis";
taxr:bioGeoStatusIn [
  rdfs:label "Guadeloupe";
  dct:spatial <http://sws.geonames.org/3579143/>;
  dwc:locationId "TDWG:LEE-GU; WOEID:23424831";
  dwc:occurrenceStatus <http://inpn.mnhn.fr/taxref/bioGeoStat#P> ] ;
taxr:bioGeoStatusIn [
  rdfs:label "New Caledonia";
  dct:spatial <http://sws.geonames.org/2139685/> ;
  dwc:locationId "TDWG:NWC-00; WOEID:23424903" ;
  dwc:occurrenceStatus <http://inpn.mnhn.fr/taxref/bioGeoStat#B> ].

<http://inpn.mnhn.fr/espece/cd_nom/60878> a skx:Label;
skx:literalForm "Delphinus delphis";
tc:authority "Linnaeus, 1758";
taxr:isPrefLabelOf <http://inpn.mnhn.fr/taxref/8.0/taxon/60878>;
taxr:vernacularName "Short-beaked common dolphin"@en.

<http://inpn.mnhn.fr/espece/cd_nom/60881> a skx:Label;
skx:literalForm "Delphinus tropicalis".
tc:authority "Van Bree, 1971";
taxr:isAltLabelOf <http://inpn.mnhn.fr/taxref/8.0/taxon/60878>;
taxr:vernacularName "Short-beaked common dolphin"@en.

```

Listing 1.6. Example representation of TAXREF entries in SKOS

be generated by a template, instead we have to write a triples map whom query filters only taxa with habitat '1'. Similarly, we must write one triples map for each habitat value, that is 8 triples maps. The same situation is observed for the 48 taxonomical ranks and 30 biogeographical statuses, that all result in specific dedicated triples maps. Consequently, many queries have to be run, some of them returning tens or hundreds of thousands of JSON documents. The same JSON documents are retrieved and parsed several times, but, each time, for the generation of triples with different properties. An analysis of the execution traces shows that, out of the 452.106 unique documents in the database, approximately 1.6 million documents are actually retrieved and processed during the translation, that is an approximate average of 4.600 documents treated per minute.

7 Discussion

xR2RML relies on the assumption that databases to translate into RDF provide a declarative query language, such that queries can be expressed directly in a mapping description. This complies with the equivalent assumption of R2RML that all RDBs support ANSI SQL. This is somehow restrictive since some NoSQL key-value stores, like DynamoDB and Riak, have no declarative query language, instead they provide APIs for usual programming languages to describe queries in an imperative manner. For xR2RML to work with those systems, a query language should be figured out along with a compiler that transforms queries into imperative code. Interestingly, this is already the case of some systems supporting the MapReduce programming model. MapReduce is conventionally supported through APIs for programming languages, however more and more systems now propose an SQL or SQL-like query language on top of a MapReduce framework (e.g. Apache Hive). Queries are compiled into MapReduce jobs. This approach is often referred to as SQL-on-Hadoop [27].

To achieve the targeted flexibility, xR2RML comes with features that are applicable independently of the type of database used. Yet, all features should probably not be applied with all kinds of database. For instance, join conditions entail joint queries. Whereas RDBs are optimized to support joins very efficiently, it is not recommended to make cross-references within NoSQL document or extensible column stores, as this may lead to poor performances. Similarly, translating a JSON element into an RDF collection is quite straightforward, but translating the result of an SQL joint query into an RDF collection is likely to be quite inefficient. In other words, because the language makes a mapping possible does not mean that it should be applied regardless of the context (database type, data model, query capabilities). Consequently, mapping designers should be aware of how databases work in order to write efficient mappings of big databases to RDF.

Like R2RML, xR2RML assumes that well-defined domain ontologies exist beforehand, whereof classes and properties will be used to translate a data source into RDF triples. In the context of RDBs, an alternative approach, the Direct Mapping, translates relational data into RDF in a straightforward man-

ner, by converting tables to classes and columns to properties [10,28]. The direct mapping comes up with an ad-hoc ontology that reflects the relational schema. R2RML implementations often provide a tool to automatically generate an R2RML direct mapping from the relational schema (e.g. Morph-RDB [29]). The same principles could be extended to automatically generate an xR2RML mapping for other types of data source, as long as they comply with a schema: column names in CSV/TSV files and extensible column stores, XSD or DTD for XML data, JSON schema²² or a JSON-LD²³ description for JSON data. Nevertheless, such schemas do not necessarily exist, and some databases like the DynamoDB key-value store are schema-less. In such cases, automatically generating an xR2RML direct mapping should involve different methods aimed at learning the database schema from the data itself.

More generally, how to automate the generation of xR2RML mappings may become a concern to map large and/or complex schemas. There exists significant work related to schema mapping and matching [30]. For instance, Clio [31] generates a schema mapping based on the discovery of queries over the source and target schemas and a specification of their relationships. Karma [32] semi-automatically maps structured data sources to existing domain ontologies. It produces a Global-and-Local-As-View mapping that can be used to translate the data into RDF. xR2RML does not directly address the question of how mappings are written, but can be complementary of approaches like Clio and Karma. In particular, Karma authors suggest that their tool could easily export mapping rules as an R2RML mapping graph. A similar approach could be applied to discover mappings between a non-relational database and domain ontologies, and export the result as an xR2RML mapping graph.

8 Conclusion and perspectives

In this paper we presented xR2RML, a language designed to describe the mapping of various types of databases to RDF, by flexibly adapting to heterogeneous query languages and data models. We analysed data models of several modern databases as well as the format in which query results are returned, and we showed that xR2RML can translate any data element within such results into RDF, relying when necessary on existing languages such as XPath and JSON-Path. We illustrated some features of xR2RML such as the generation of RDF collections and containers, and the ability to deal with mixed data formats, e.g. when the column of a relational table stores data formatted in another syntax like XML, JSON or CSV.

Principles of the xR2RML mapping language were validated in a prototype implementation supporting several RDBs and the MongoDB NoSQL document store. The prototype was used in a real-world use case to perform the translation of a taxonomical reference of more than 450.000 taxa, represented as JSON documents stored in a MongoDB instance, into a SKOS thesaurus. The data

²² <http://json-schema.org/>

²³ <http://www.w3.org/TR/json-ld/>

materialization approach we implemented proved to be effective, although the use case underlined scaling limitations with regards to execution time span and memory consumption. In particular, this approach cannot scale to big data sets. Dealing with big data sets requires the data to remain in legacy databases, and that the translation to RDF be performed on demand through the xR2RML-based rewriting of SPARQL queries into the source database query language. In this regard, existing works related to RDBs should be leveraged [24,33].

References

1. B. He, M. Patel, Z. Zhang, and K. C.-C. Chang, "Accessing the deep web," *Communications of the ACM*, vol. 50, no. 5, pp. 94–101, 2007.
2. S. Das, S. Sundara, and R. Cyganiak, "R2RML: RDB to RDF mapping language," 2012.
3. A. Dimou, M. V. Sande, J. Slepicka, P. Szekely, E. Mannens, C. Knoblock, and R. V. d. Walle, "Mapping hierarchical sources into RDF using the RML mapping language," in *Proc. of ICSC'2014*, pp. 151–158, IEEE, 2014.
4. M. T. Roth and P. Schwartz, "Don't scrap it, wrap it! A wrapper architecture for legacy data sources," in *Proc. of VLDB'1997*, pp. 266–275, 1997.
5. J. Melton, J. E. Michels, V. Josifovski, K. Kulkarni, and P. Schwarz, "SQL/MED: a status report," *ACM SIGMOD Record*, vol. 31, no. 3, pp. 81–89, 2002.
6. A. Schwarte, P. Haase, K. Hose, R. Schenkel, and M. Schmidt, "FedX: Optimization techniques for federated query processing on linked data," in *Proc. of ISWC'11*, pp. 601–616, 2011.
7. M. Acosta, M. Vidal, T. Lampo, J. Castillo, and E. Ruckhaus, "ANAPSID: an adaptive query processing engine for SPARQL endpoints," in *Proc. of ISWC'11*, pp. 18–34, 2011.
8. A. Gaignard, "Distributed knowledge sharing and production through collaborative e-science platforms. PhD thesis," 2013.
9. D.-E. Spanos, P. Stavrou, and N. Mitrou, "Bringing relational databases into the semantic web: A survey," *Semantic Web Journal*, vol. 3, no. 2, pp. 169–209, 2012.
10. J. Sequeda, S. H. Tirmizi, s. Corcho, and D. P. Miranker, "Survey of directly mapping SQL databases to the semantic web," *Knowledge Eng. Review*, vol. 26, no. 4, pp. 445–486, 2011.
11. F. Michel, J. Montagnat, and C. Faron-Zucker, "A survey of RDB to RDF translation approaches and tools," 2014. Research report. ISRN I3S/RR 2013-04-FR.
12. S. Bischof, S. Decker, T. Krennwallner, N. Lopes, and A. Polleres, "Mapping between RDF and XML with XSPARQL," *Journal on Data Semantics*, vol. 1, no. 3, pp. 147–185, 2012.
13. P. Fennell, "Schematron - more useful than you'd thought," in *Proc. of the XML London 2014 Conference*, pp. 103–112, 2014.
14. F. Breitling, "A standard transformation from XML to RDF via XSLT," *Astrophysical Journal*, vol. 330, p. 755, 2009.
15. N. Bikakis, C. Tsinaraki, I. Stavrakantonakis, N. Gioldasis, and S. Christodoulakis, "The SPARQL2XQuery interoperability framework," *CoRR*, vol. abs/1311.0536, 2013.
16. A. Langegger and W. Wöss, "XLWrap - querying and integrating arbitrary spreadsheets with SPARQL," in *Proc. of ISWC'2009*, 2009.

17. F. Scharffe, G. Ateazing, R. Troncy, F. Gandon, S. Villata, B. Bucher, F. Hamdi, L. Bihanic, G. Képéklian, F. Cotton, and others, “Enabling linked data publication with the Datalift platform,” in *Proc. of the AAAI workshop on semantic cities*, 2012.
18. A. Dimou, M. Vander Sande, P. Colpaert, R. Verborgh, E. Mannens, and R. Van de Walle, “RML: A generic language for integrated RDF mappings of heterogeneous data,” in *Proc. of the 7th LDOW workshop*, 2014.
19. R. Hecht and S. Jablonski, “NoSQL evaluation: A use case oriented survey,” in *Proc. of CSC’2011*, pp. 336–341, IEEE Computer Society, 2011.
20. S. K. Gajendran, “A survey on NoSQL databases (technical report),” 2013.
21. K. W. Ong, Y. Papakonstantinou, and R. Vernoux, “The SQL++ unifying semi-structured query language, and an expressiveness benchmark of SQL-on-Hadoop, NoSQL and NewSQL databases (submitted),” *CoRR*, vol. abs/1405.3631, 2014.
22. B. Kolev, P. Valduriez, R. Jimenez-Peris, N. Martínez-Bazan, and J. Pereira, “CloudMdsQL: Querying heterogeneous cloud data stores with a common language,” in *Proc. of the BDA’2014 Conference*, 2014.
23. F. Michel, L. Djimenou, C. Faron-Zucker, and J. Montagnat, “xR2RML: Relational and non-relational databases to RDF mapping language,” 2014. Research report. ISRN I3S/RR 2014-04-FR v3.
24. F. Priyatna, O. Corcho, and J. Sequeda, “Formalisation and experiences of R2RML-based SPARQL to SQL query translation using Morph,” in *Proc. of WWW’2014.*, 2014.
25. P. Gargominy, S. Tercerie, C. Régnier, T. Ramage, C. Schoelinck, P. Dupont, E. Vandell, P. Daszkiewicz, and L. Poncet, “TAXREF v8.0, référentiel taxonomique pour la France: Méthodologie, mise en oeuvre et diffusion,” in *Rapport SPN 2014 - 42*, 2014.
26. C. Callou, F. Michel, C. Faron-Zucker, C. Martin, and J. Montagnat, “Towards a Shared Reference Thesaurus for Studies on History of Zoology, Archaeozoology and Conservation Biology,” in *ESCW 2015, workshop Semantic Web For Scientific Heritage (SW4SH)*, (Portoroz, Slovenia), 2015.
27. A. Floratou, U. F. Minhas, and F. Ozcan, “Sql-on-hadoop: Full circle back to shared-nothing database architectures,” *Proc. of the VLDB Endowment*, vol. 7, no. 12, 2014.
28. M. Arenas, A. Bertails, E. Prud’hommeaux, and J. Sequeda, “A direct mapping of relational data to RDF,” 2012.
29. L. F. de Medeiros, F. Priyatna, and O. Corcho, “MIRROR: Automatic R2RML mapping generation from relational databases,” in *Submission to ICWE 2015*, 2015.
30. P. Shvaiko and J. Euzenat, “A survey of schema-based matching approaches,” in *Journal on Data Semantics IV*, pp. 146–171, Springer, 2005.
31. R. Fagin, L. M. Haas, M. Hernandez, R. J. Miller, L. Popa, and Y. Velegrakis, “Clio: Schema mapping creation and data exchange,” in *Conceptual Modeling: Foundations and Applications*, pp. 198–236, Springer, 2009.
32. C. A. Knoblock, P. Szekely, J. L. Ambite, A. Goel, S. Gupta, K. Lerman, M. Muslea, M. Taherian, and P. Mallick, “Semi-automatically mapping structured sources into the semantic web,” in *Proc. of ESWC’2012*, pp. 375–390, Springer, 2012.
33. J. F. Sequeda and D. P. Miranker, “Ultrawrap: SPARQL execution on relational data,” *Web Semantics: Science, Services and Agents on the WWW*, vol. 22, pp. 19–39, 2013.