Richard Cyganiak
richard@cyganiak.de

Christian Bizer
chris@bizer.de

Freie Universität Berlin

D2R Server: A Semantic Web Front-end to Existing Relational Databases

RDF and the SPARQL query language and protocol provide standardized ways of exposing and linking data sources on the Web.

D2R Server is a system for publishing content from relational databases on the Semantic Web. By rewriting requests to SQL, D2R Server can run on top of live databases without having to replicate its content into a native RDF store.
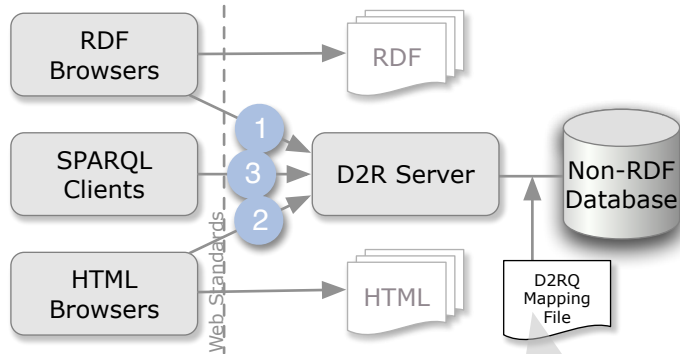
Open Data and the Semantic Web

Public and private organizations can benefit from **making data available** in machine-readable formats. Such data is most valuable when it is **easily accessible** and can be **re-used** by integrating it with other data. Semantic Web technologies support this through

- built-in globally unique identifiers (URIs),
- powerful options for linking and integrating data published by different parties,
- the ability to mix and match vocabularies for the description of a single resource.

Data published on the Semantic Web can be queried using the **SPARQL** query language, can be navigated with **RDF browsers** like Tabulator and Piggy Bank, and is accessible to RDF-consuming **Web crawlers**.

1 RDF Representations

The mapping file assigns URIs to objects in the database. HTTP Requests on these URIs are answered with an RDF description of the object.

By requesting the URIs of other resources mentioned in the description, RDF browsers and crawlers can navigate the Web of data. Descriptions can refer to external RDF resources published by third parties.

2 HTML Representations
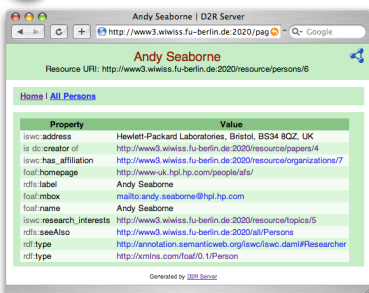
D2RQ Mapping File

Describes how to translate between application-specific database schemas and RDFS schemas or OWL ontologies.

Part of a database schema and a corresponding mapping file in N3:


```
map:movies a d2rq:ClassMap;
    d2rq:uriPattern "movies/@@movies.ID@@";
    d2rq:class imdb:Movie;
    d2rq:condition "movies.status='published'";
    .
map:title a d2rq:PropertyBridge;
    d2rq:belongsToClassMap map:movies;
    d2rq:property imdb:title;
    d2rq:column "movies.title";
    .
map:genre a d2rq:PropertyBridge;
    d2rq:belongsToClassMap map:movies;
    d2rq:property imdb:genre;
    d2rq:refersToClassMap map:genres;
    d2rq:join
        "movies.ID = movies2genres.movieID";
    d2rq:join
        "genres.ID = movies2genres.genreID";
    .
map:genres a d2rq:ClassMap;
    d2rq:uriPattern
        "http://genres/@@genres.genre@@";
    d2rq:class imdb:Genre;
    .
```

3 SPARQL Queries

Queries and results are transmitted using the SPARQL protocol. The server breaks queries into simpler parts and rewrites them on-the-fly to SQL for processing in the database.

```
SELECT ?title WHERE {
    ?movie :title ?title ;
        :year "1977" ;
        :genre ?genre .
    ?genre rdfs:label "Sci-Fi" . }
```

```
SELECT DISTINCT movies.title
  FROM movies, movies2genres, genres
 WHERE movies.year=1977
   AND genres.ID=movies2genres.genreID
   AND movies.ID=movies2genres.movieID
   AND genres.genre='Sci-Fi';
```