

ONTOLOGY-BASED DATA ACCESS AND VISUALIZATION OF BIG VECTOR AND RASTER DATA

Konstantina Bereta, George Stamoulis and Manolis Koubarakis

National and Kapodistrian University of Athens

ABSTRACT

We present an Ontology-based data access technique for the integration of geospatial data from different sources and different formats using ontologies and mappings. Our technique is implemented in the system Ontop-spatial, which enables the execution of GeoSPARQL queries over geospatial relational databases storing vector or raster data by performing on-the-fly GeoSPARQL-to-SQL translation using ontologies and mappings. Ontop-spatial is generally more efficient than traditional geospatial RDF stores. Linked geospatial data can then be visualised on the map using Sextant, an online tool for browsing and visualisation of linked geospatial data.

Index Terms— GeoSPARQL, Ontology-based Data Access, Raster data, Visualisation

1. INTRODUCTION

Previous projects TELEIOS, LEO and Melodies funded by FP7 ICT, and OBEOS funded by ESA have demonstrated the use of linked data in Earth Observation (EO). The current H2020 project Copernicus App Lab (<http://www.app-lab.eu/>) goes one step further by making data from three Copernicus services (Land, Marine and Atmosphere) available on the Web as linked data to aid their utilization by mobile developers. In previous projects, it has been assumed that EO data are transformed from their original formats (shapefiles, spatially-enabled relational databases, GeoTIFF, NetCDF etc.) into RDF, stored in geospatial RDF stores and queried using geospatial extensions of SPARQL to develop interesting applications. In this paper, we present the system Ontop-spatial which enables the creation of *virtual RDF graphs over EO data stored in their original formats* using ontologies and mappings. Ontop-spatial allows EO data centers to make their data available as linked data that can be queried using the OGC standard GeoSPARQL [1], without first having to translate this data into RDF. Ontop-spatial scales to big geospatial data and it is more efficient than related geospatial RDF stores. Ontop-spatial is available as open source at <https://ontop-spatial.di.uoa.gr>.

This work has been funded by the EU project Copernicus App Lab (730124).

Ontop-spatial adopts the Ontology-Based Data Access (OBDA) paradigm pioneered by the Semantic Web community, and it is the *first geospatial OBDA system*. Ontop-spatial is able to connect to geospatial databases and create geospatial RDF graphs on top of them, using ontologies (that are extensions of the GeoSPARQL ontology) and R2RML mappings. Figure 1 shows graphically the classes of the GeoSPARQL ontology.




Fig. 1. Classes of the GeoSPARQL ontology

R2RML (<https://www.w3.org/TR/r2rml/>) is the standard language for encoding how relational data is mapped into RDF terms. This virtual approach avoids the need of materialization and facilitates data integration, as it enables users to pose the same GeoSPARQL queries they would pose over the materialized RDF data. GeoSPARQL queries are translated by Ontop-spatial on-the-fly into the respective SQL queries with spatial operators, and are evaluated in the geospatial DBMS. Currently, PostGIS, Spatialite and Oracle Spatial are supported as back-ends. The first version of Ontop-spatial dealing with vector data only has been presented in [2]. This version has been used in three environmental applications in the context of project MELODIES, and in a marine-security application in the context of German national project EMSec.

The contribution of our approach with respect to the Big Data from space dimensions are the following. **Volume:** Ontop-spatial outperforms the state-of-the-art in geospatial RDF stores and is able to process tens of Gigabytes of data containing complicated geometries. **Velocity:** When data gets frequently updated, using traditional triple stores is inefficient, as batches of data need to be converted and materialized as RDF triples each time they arrive. Our approach eliminates as much as possible the need for materializing data and it is suitable for data sources that get frequently updated

(e.g., streams). **Variety:** With raster and OPeNDAP support in place, Ontop-spatial becomes the first GeoSPARQL query engine that is able to process such a wide variety of geospatial formats, enabling geospatial data integration using ontologies and mappings. **Value:** Exposing geospatial data as virtual RDF triples that can be accessed in the Web through standard (Geo)SPARQL endpoints enables the interlinking of EO data with other data (e.g., open data) increasing their value, as data from multiple geospatial sources can be combined and rich queries can be expressed over them.

Since vector and raster data is made available as linked geospatial data, we can visualise them using the tool Sextant [3]. Sextant is a web tool for browsing and visualising linked geospatial and temporal data and is able to connect to different (Geo)SPARQL endpoints and project geospatial results on a map. Each layer of the map can be created either by retrieving results available on the web via (Geo)SPARQL endpoints or by uploading geospatial files (e.g., KML, GEOTIFF, etc.) directly. Sextant combines the functionality offered by traditional GIS with the ability to combine data that is stored locally and data that is already available on the web, to create maps as web resources enabling collaborative on-line editing.

2. ONTOLOGY-BASED DATA ACCESS ON TOP OF VECTOR AND RASTER DATA

The work described in [2] documented in detail how Ontop-spatial can be used to create virtual geospatial RDF graphs on top of geospatial DBMSs. Although there are many geospatial RDF stores implementing -usually part of- GeoSPARQL, there was no OBDA system with GeoSPARQL support until the creation of Ontop-spatial. Following to this work, we have now extended Ontop-spatial with the ability to query raster data as well. Querying raster data sources using declarative query languages can also be done using array DBMSs such as Rasdaman, MonetDB and SciDB. As GeoSPARQL does not include support for raster data, in our approach we do not deviate from the standard but instead: i) we overload existing vector GeoSPARQL operators such as `geof:sfIntersects` to be used with raster data as well, and ii) in the mappings, we use the raster functions supported by the underlying DBMS (e.g., PostGIS with the raster support).

More recently, work on the SciSPARQL query language showed how to query grid coverages using a hybrid data store composed of Rasdaman and a main-memory RDF store [4]. We deviate from this approach by proposing the encapsulation of raster data functions in the mappings instead of extending (Geo)SPARQL with array functionalities, so that not every raster cell needs to be represented in RDF.

The problem of representing and querying raster data as linked data has also been discussed in the recent working note "Coverages in Linked Data" by the OGC/W3C Spatial Data on the Web working group

(https://www.w3.org/2015/spatial/wiki/Coverages_in_Linked_Data).

None of the geospatial extensions of the framework of RDF and SPARQL, such as stRDF and stSPARQL and GeoSPARQL have considered support for raster data. The main challenge that lies behind this is twofold. First, a raster file is associated with a geometry only as a whole. It is not straightforward to associate separate raster cells to a geometry; they have to be vectorized first (i.e., translated into polygons). Second, every raster cell is associated with one or more values. In order to convert all information contained in a raster file into RDF, then multiple triples should describe a raster cell, producing a large amount of triples for a whole raster file. However, not all of this information is needed. In most of the use cases, only the information that derives from a raster file and satisfies certain criteria (e.g., value constraints) is all that is needed to be converted into RDF. This means that the raster file needs to be processed and then the results of this processing are useful as RDF, while any other information is redundant. These challenges have discouraged the scientific community from converting and materializing raster data to RDF. The following example describes how raster data can be mapped into virtual RDF data. For the convenience of the reader, we present the mappings using the OBDA native language of Ontop instead of R2RML, as it is more compact and readable, but R2RML is also supported in the system.

```
mappingId chicago2
target geo:{geom} rdf:type f:rastCell;
        geo:asWKT {geom} .
source  select ST_DumpAsPolygons(rast)
as geom from chicago;
```

In the example described above, a GeoTIFF image has been imported into a PostGIS database as relation `chicago`. The mapping shows how raster data stored in column `rast` are mapped to geometries in WKT format, after they are vectorized, using the PostGIS `ST_DumpAsPolygons` function. This is a procedure that allows domain experts to use all geometries that they may have in a database uniformly, and execute spatial operations involving vector and raster geometries. Domain experts usually perform this vectorization step as part of pre-processing. In the mapping described above, we show how this can be done on-the-fly, using Ontop-spatial.

In the project Copernicus App Lab, Ontop-spatial has also been extended to support data sources made available via OPeNDAP services offered by our partner Dutch company RAMANI. OPeNDAP is a framework for accessing scientific data (<https://www.opendap.org/>) which is widely used by Earth scientists, as it is popular in large organizations such as NASA and NOAA. Earth science data can be consumed by using a specific OPeNDAP client. To make data provided by OPeNDAP services available as linked data, the data should be downloaded, materialized and then converted into RDF using custom programs, as existing applications that convert geospatial data into RDF do not offer support for OPeNDAP.

The approach that we describe in this paper enables the creation of virtual geospatial RDF graphs on top of data that is accessible through OPeNDAP on-the-fly, without materializing the original data or the RDF data.

Ontop-spatial has been extended with an adapter that enables it to retrieve data from an OPeNDAP server, create a table view on-the-fly, populate it with this data and create virtual semantic geospatial graphs over it. To achieve this, Ontop-spatial utilizes the system MadIS (<https://github.com/madgik/madis>) as a back-end. MadIS is an extensible relational database system built on top of the APSW SQLite wrapper. MadIS is a framework that provides a Python interface so that users can easily implement user-defined functions (UDFs) as row, aggregate functions, or virtual tables. We used MadIS in order to create a new UDF, named `Opendap`, that is able to create and populate a virtual table on-the-fly with data retrieved from an OPeNDAP server. In this way, Ontop-spatial enables users to pose GeoSPARQL queries on top of OPeNDAP data sources without materializing any triples or tables.

An example is provided below.

```
mappingId opendap_mapping
target lai:{id} rdf:type lai:Observation ;
      lai:{id} lai:hasLai {LAI}^^xsd:float;
      lai:detectionTime {time}^^xsd:dateTime;
      geosparql:asWKT {wkt}^^geo:wktLiteral .
source select id, LAI, time, wkt
      from (ordered opendap
            url:https://analytics.ramani.ujuizi.com/
            %28https://ramani.ujuizi.com/
            thredds/dodsC/Copernicus-Land-timeseries-
            global-LAI%29/readdods/.LAI/)
            where LAI > 0
```

In this mapping, the `source` is a Leaf Area Index (LAI) dataset with resolution of 100 meters is provided through an OPeNDAP server. The dataset contains observations about the LAI values of areas, as well as the time and location for each observation. The MadIS operator `Opendap` retrieves this data and populates a virtual database table with the schema `(id, LAI, time, wkt)`. The column `id` was not originally in the dataset but it is constructed from the location and time when the observation is taken. The LAI column stores the LAI values of an observation as `float` values. The attribute `time` represents the timestamp of an observation in `datetime` format. In the original dataset temporal values are represented as numeric values. The meaning of these values is described in the metadata. For example, it can be days or months since a fixed timestamp. Unfortunately, this is not a standard representation that we would have available if we had imported the dataset into a geospatial database. Because of the fact that the `Opendap` operator is implemented as an SQL user-defined operator, it can be embedded into any SQL query. So we refined the data that we want to be translated into virtual RDF terms by adding an SQL filter to the query to eliminate the negative or zero LAI values by filtering them out at an intermediate level, so that i) we do not change the values of the original dataset and ii) we provide only the cor-

rect values to the users so that they do not need to handle the noise themselves (e.g., by using GeoSPARQL filters or custom code).

The `target` part of the mapping encodes how the relational data can now be mapped into RDF terms. Every row in the virtual table describes an instance of the class `lai:Observation`. The values of the LAI column populate the triples that describe the LAI values of the Observation, and the values of the columns `time` and `wkt` populate the triples that describe the time and location of the observations accordingly.

Given the mappings provided above, we can pose the the following GeoSPARQL query to retrieve the Leaf Area Index values and the geometries of areas

```
select distinct ?s ?g ?lai where {
  ?s lai:hasLai ?lai .
  ?s geo:asWKT ?g }
```

Notably, both translation steps are performed on-the-fly and only after a GeoSPARQL query is posed to the system. This approach goes considerably beyond the previous version of Ontop-spatial that could only connect to an existing database with materialized tables, as well as the default, non-spatial version of Ontop and any other RDB2RDF system. OBDA systems traditionally connect to an existing database with materialized tables and access it before a query is fired in order to collect metadata, etc. The exact schema of the database tables is known beforehand. The approach that we propose in this paper is *schema-agnostic*: Ontop-spatial does not know the schema of the data as there is no database materialized. The virtual table is only created on-the-fly.

Ontop-spatial can be available as a GeoSPARQL endpoint, and thus can be used both by federation and interlinking engines. For example, one can use the interlinking tool Silk (<http://silkframework.org/>) to interlink Copernicus data that is accessible as RDF graphs using Ontop-spatial with linked data that is available using standard (Geo)SPARQL endpoints.

We have evaluated Ontop-spatial by extending the benchmark Geographica with support for the evaluation of OBDA systems. Geographica (<http://geographica.di.uoa.gr>) was initially designed to evaluate the performance of geospatial RDF stores. We compared Ontop-spatial with the state-of-the-art geospatial RDF store Strabon (<http://strabon.di.uoa.gr>) [5]. Strabon has also been developed by our group and has been shown to be the most efficient geospatial RDF store available today [6]. Our evaluation showed that Ontop-spatial generally achieves *significantly better performance* than Strabon, often by orders of magnitude, when a large number of geospatial intermediate results are generated during the evaluation of a query. For example, Ontop-spatial is able to execute spatial selections and spatial joins against a 30 GB dataset that contains complex geometries (i.e., from points to polygons con-

taining thousands of points) in less than a second. More details about are work in this respect are provided in [7].

3. VISUALISATION OF LINKED VECTOR AND RASTER DATA

In order to visualise and explore Linked Vector and Raster data we use the tool Sextant [3]. Sextant (<http://sextant.di.uoa.gr/>) is a web based platform for visualizing, exploring and interacting with linked geospatial and temporal data.

While Sextant offers a variety of features for domain experts, it provides a user-friendly interface and some complementary features such as predefined queries and explore functionality, that enable non-experts to take advantage of semantic web technologies, and share the benefits of the linked open geospatial Web.

The core feature of Sextant is the ability to create thematic maps by combining geospatial and temporal information that exists in a number of heterogeneous data sources ranging from standard SPARQL endpoints, to SPARQL endpoints following the standard GeoSPARQL defined by the Open Geospatial Consortium (OGC), or well-adopted geospatial file formats, like KML, GML and GeoTIFF.

These maps are modeled using an ontology, so they can be expressed in RDF and then stored, creating a map registry. In Sextant we use the spatiotemporal RDF store Strabon to create such registries. Each registry is also a resource and can be accessed with a unique URL that the user can provide in the User Interface (UI) in order to connect to the underlying RDF store and use it to store, load or search for maps. Map registries as a result play an important role in sharing and collaborative viewing of maps.

Figure 2 depicts a map in Sextant that was created by combining information from various vector and raster datasets.

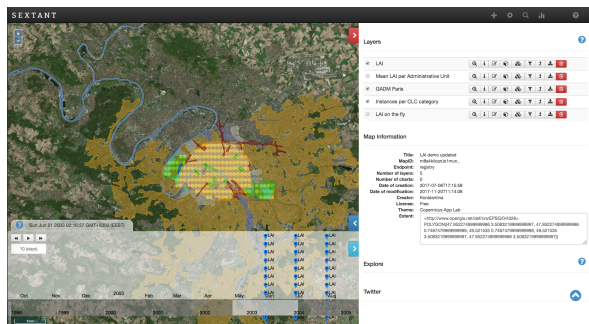


Fig. 2. Leaf Area Index Demo map

This demo map consists of five layers that demonstrate the Paris Green Capital use case of the Copernicus App Lab project:

- Global Administrative Divisions of Paris

- Leaf Area Index in Paris
- Instances per Corine Land Cover Category
- Mean Leaf Area Index per Administrative Division
- OpenStreetMap Parks in Paris

Each layer in this demo map is produced using a GeoSPARQL query, since Sextant can detect and visualise on the world map the spatial and temporal dimension from a query result set.

4. REFERENCES

- [1] “Open Geospatial Consortium. OGC GeoSPARQL - A geographic query language for RDF data,” OGC Candidate Implementation Standard, 2012.
- [2] K. Bereta and M. Koubarakis, “Ontop of geospatial databases,” in *ISWC 2016*.
- [3] Charalampos Nikolaou, Kallirroï Dogani, Konstantina Bereta, George Garbis, Manos Karpathiotakis, Kostis Kyzirakos, and Manolis Koubarakis, “Sextant: Visualizing time-evolving linked geospatial data,” *Web Semantics: Science, Services and Agents on the World Wide Web*, vol. 35, no. 1, 2015.
- [4] Andrej Andrejev, Dimitar Misev, Peter Baumann, and Tore Risch, “Spatio-temporal gridded data processing on the semantic web,” in *DSDis 2015*.
- [5] Kostis Kyzirakos, Manos Karpathiotakis, and Manolis Koubarakis, “Strabon: A Semantic Geospatial DBMS,” in *ISWC*, Philippe Cudr-Mauroux and et al., Eds. 2012, vol. 7649 of *LNCS*, pp. 295–311, Springer.
- [6] G. Garbis, K. Kyzirakos, and M. Koubarakis, “Geographica: A Benchmark for Geospatial RDF Stores,” *ISWC 2013*.
- [7] Konstantina Bereta and Manolis Koubarakis, “Creating Virtual Semantic Graphs on Top of Big Data from Space,” in *Proc. of the 2017 Conference on Big Data from Space (BiDS)*, 2017.