

Using RDF to Query Spatial XML

Jose Eduardo Córcoles and Pascual González

LOuISE Research Group. Dept. Informática.
Universidad de Castilla-La Mancha
02071, Albacete (Spain)
{corcoles, pgonzalez}@info-ab.uclm.es

Abstract. In tackling the *Semantic Web*, a rich domain that requires special attention is the Geospatial Semantic Web. A contribution to the Geospatial semantic Web is the definition of an approach for integrating non-spatial resources like HTML, PDF, GIF, etc., with spatial XML resources represented by Spatial XML [4]. However, in order to put this approach into practice it is necessary to solve the problem of developing an integration system for querying spatial XML resources stored in different sources. In this paper, we have implemented an approach for querying spatial and non-spatial information represented in the Geographical Markup Language (GML). The approach uses RDF to integrate the spatial XML documents. A performance study has been carried out and the results are given.

1 Introduction

A rich domain that requires special attention is the Geospatial Semantic Web [1]. The enormous variety of encoding of geospatial semantics makes it particularly challenging to process requests for geospatial information. In the future, the Geospatial Semantic Web will allow the returning of both spatial and non-spatial resources to simple queries, using a browser. For example, a query “*lakes in Maine*” should return all relational resources with lakes in Maine (pictures, text, ...) in different formats (XML, HTML, JPG, PDF, References, ...) [1].

However, in the same way as with the Semantic Web, in order to approach the Semantic Geospatial Web it is necessary to solve several problems. One of these is the integration of spatial and non-spatial resources with different schemas stored in different sources.

Tackling the integration of spatial information on the Web is not a simple task since there are several High level problems (e.g. heterogeneity for representing spatial information, for using spatial operators, for defining the semantic of the objects, etc. [2]) and Low level problems (e.g. the sources may store large amounts of incomplete spatial data, which may make it necessary to join the results of queries with spatial joins [3]).

In order to contribute to the Geospatial Semantic Web, in this paper we study a prototype for querying spatial XML resources. The main task of this approach is to provide users with a unique interface for querying spatial XML resources with different schemas, independently of their actual organization and location. It provides

the infrastructure for formulating structured spatial queries by taking into consideration the conceptual representation of a specific domain in the form of an ontology. The resources are integrated using RDF. This work is based on [4], and we detail how to use RDF for integrating non-spatial resources (HTML, PDF, etc.) with spatial XML resources represented by GML. The most novel and critical feature of this approach is the querying of spatial XML resources, because it uses a different way from that of querying and relating non-spatial resources. Therefore, our prototype is focused on this alone. After proving the viability of this prototype, making a prototype in accordance with [4] is a simply task.

In our study the spatial information is represented by GML because it is an XML encoding for the transport and storage of spatial/geographic information, including both spatial features and non-spatial features. The mechanisms and syntax that GML uses to encode spatial information in XML are defined in the specification of OpenGIS [5]. Thus, GML allows a more homogeneous and flexible representation of the spatial information.

Query mediation has been extensively studied in the literature. Directly concerned with the integration of XML resources, it is worth noting C-Web Portal [6] and [7]. C-Web Portal supports the integration of non-spatial resources on the Web, and it provides the infrastructure for formulating structured queries by taking into consideration the conceptual representation of a specific domain in the form of an ontology. On the other hand, [7] proposes a *mediator* architecture for the querying and integration of Web-accessible XML data resources (non-spatial data). Its contribution is the definition of a simple but expressive mapping language, following a *local as view* approach and describing XML resources as local views of some global schema. Both of these approaches have aims in common with our approach, but they are only focused on non-spatial information.

In relation to spatial XML integration, the approaches developed by [8], [9] and [10] stand out. [8] presents a mediation system that addresses the integration of GIS data and tools, following a *global-as-view* approach. It has a multi-tier client-server architecture based on WFS and uses standard wrappers to access data, extended by derived wrappers that capture additional query capabilities. [9] extends the MIX wrapper-mediator architecture for integrating information from spatial information systems and searchable databases of geo-referenced imagery. MIX is focused on integrating geo-referenced imagery but our approach is focused on spatial geometries. On the other hand, [10] designed a novel approach for integrating GML resources. The proposed architecture uses a Catalog expressed by RDF to relate the GML resources. Although [10] has the same aims as [4], it has a different focus.

An overview of the architecture offered in this work is detailed in Section 2, which contains the most important features of our prototype. Section 3 provides some results from our performance studies. Section 4 contains conclusions and projected future work.

2 Overview

The main task of an integration mediator is to provide the users with a unique interface for querying the data, independently of its actual organization and location [11]. This interface, or global schema, is described as an *ontology* expressed with

RDF(S). As used here, an *ontology* denotes a light-weight conceptual model and not a hierarchy of terms or a hierarchy of concepts [7]. The global schema can be viewed as a simple object-oriented data model. Hence, a global schema can be viewed as defining a database of objects, connected by roles, with the concept extents related by subset relationships as per the *isA* links in the schema. Since it is an integration schema, this is a *virtual* database. The actual materialization exists in the resources.

Users pose queries in terms of global schema (RDFS). As a consequence, the mediator system must contain a module (*Solve mapping*) that uses the resource descriptions in order to translate a user query into a query that refers directly to the schemas of the GML resources. For this purpose, we establish a correspondence between each resource and the global schema using RDF instances. In the *Solve mapping* module, an execution plan is not necessary because a query will be executed over a source if it fully satisfies all attributes of the query.

In addition, this approach supplies *descriptions* of the resources and specifies constraints on resource contents (e.g. contains information about “Madrid”). These *descriptions* are stored as alphanumeric data in RDF instances (Filter Resources).

A Wrapper receives a query over a schema of the resources from the mediator system. Here, the query is executed and the results of the query are expressed in GML format, which are then returned to the mediator.

This process is a well-known pattern in applications over non-spatial XML documents. However, to apply this architecture efficiently over spatial XML documents it is necessary to take into account the following points:

1. The spatial query language used for the users should have the same semantic as the query language used in the sources, so the translation between both languages is easier. Because of this, we have used the same spatial language to query the ontology and the GML data model on the sources[4]. This query language was born as a spatial query language over spatial semi-structured XML documents. The data model and the algebra underlying the query language are defined in [12]. The query language has a familiar *select-from-where* syntax and is based on SQL. It includes a set of spatial operators (disjoint, touches, etc.), and includes traditional operators (=, >, <, ...) for non-spatial information.
2. Unlike XML, it is very difficult to make efficient queries over GML because the spatial operator requires a spatial index that cannot be created directly over GML. Due to this an efficient method for storing GML documents is necessary. In [13] we studied three storage models for storing and retrieving GML documents over RDBMS. This work concluded that the LegoDB approach [14] obtains the best performance. We have therefore used this approach in our implementation.
3. Obviously, in schema integration the mediator can know when a resource has schema to satisfy a query. However, only when the query is executed in the resources does the mediator know whether the information stored in the resource satisfies the *where clause*. Thus, for example, it is possible that 100 resources have schema to satisfy the query, but only 10 have the information that the user wants. In this case, posting the query to each resource (100) is inefficient. To solve this, our approach uses the above-mentioned *description* to obtain *a priori* a set of candidate resources (in the same way as a *Catalog* in [15]).

In the following section the most important features of our prototype are detailed.

2.1 Ontology

The solution proposed in this paper is based on [4], which details how to use RDF(S) for querying spatial XML resources (GML). [4] is based on *Community Web Portal (C-Web)* [16]. *C-Web* essentially provides the means to select, classify and access, in a semantically meaningful and ubiquitous way, various information resources for diverse target audiences. However, in [4] we do not only use the RDF(S) to relate resources (html, pdf, jpg,...) situated in different web sites; we also use it to know how a schema satisfies an ontology. These schemas are always expressed as DTDs for each source. Each site can offer the schemas that satisfy, fully or partially, the ontology defined for an interest community. Therefore, unlike the original C-Web, it is possible to apply spatial operators (comparatives: cross, overlap, touch; analysis: Area, Length) over the resources provided that they represent geometry information with GML.

In Figure 1an example of the *Portal Schema* and its instances is shown. Due to RDF's capability for adding new feature and geometry types in a clear and formal manner, this example has been carried out extending the geospatial ontology defined by OpenGIS [5], where the class (*Geometry, LineString, etc*) and properties (*coordinates, PolygonMember, etc*) are defined. The example shows a simplification of a City Model. In this, *CityModel* contains *Blocks* and a *Block* contains *Parcels*, and a *CityModel* has the properties *name, population* and *Category*.

The most important feature of this approach is that there are two kinds of properties: (1) properties that describe the data offered by the resources (*descriptions*), represented by a rectangle filled with white; (2) properties that represent the name of the attributes of a schema (DTD) that have the same semantic, represented by a rectangle filled with gray.

The first set of properties such as *Name* and *Category* in *CityModel* or *Type* in *Parcel* make the *Catalog*. It allows a spatial query to be refined before it is posted to the sources. For example, if we wanted to execute a query about parcels in "Albacete", it would not be necessary to execute the query in all the resources with information about parcels, but only in those resources related to the *CityModel* called "Albacete". At the beginning, users filter resources in accordance with their description in order to get the candidate resources, and then users can apply a spatial query over these alone.

Evidently, a large *Catalog* allows better filtering of a query, but it leads to difficult management. Owing to this, the size of the *Catalog* should be defined by consensus for each problem.

With regard to the second set of properties, in our prototype the attributes are referenced by beginning with the root of the XML documents. For instance, the resource &r2 (*Block*) located in *www.cities.com/Blocks.gml* represents in a DTD the "Boundary" property with the attribute "world.city.BoundaryCity", and the same property is represented by &r3 with "Block.BoundaryBy". Figure 2 shows an example of a data model for the resource &r3. This data model is obtained to start from a GML document [12]. Note that *dot notation* is used to describe the attributes in the spatial XML data model. It is the same syntax used in our query language mentioned above. Thanks to this, a translation between notations (query and mapping) is not necessary. In addition, *wildcards* and other features related to semi-structures can be used to define attributes [12].

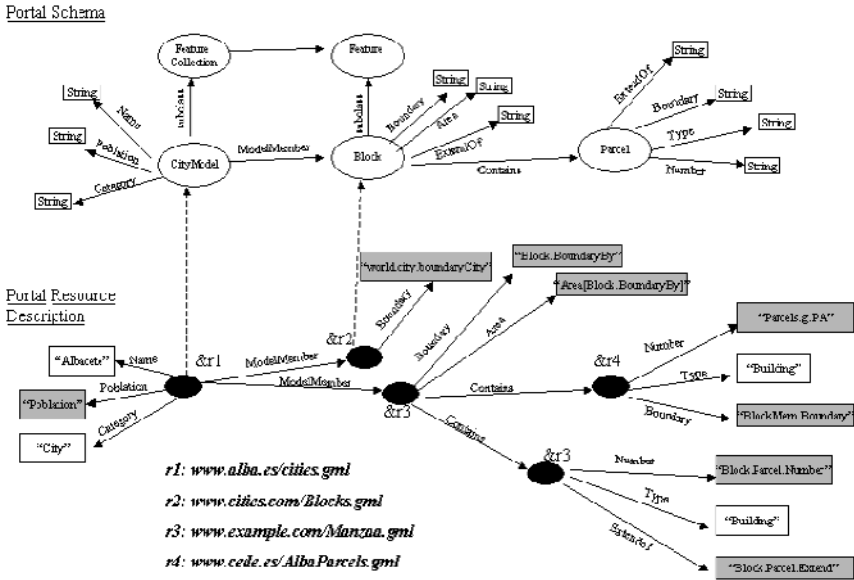


Fig. 1. Portion of Catalog of a CityModel.

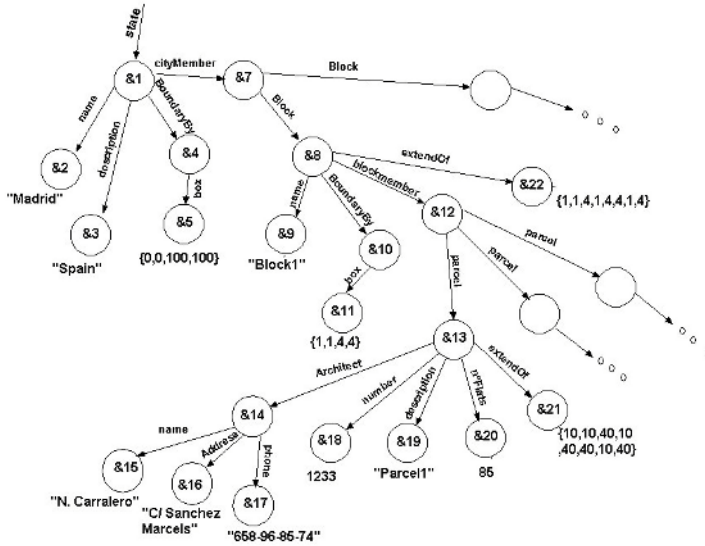
On the other hand, spatial operators defined with the syntax of our query language can also be used to enrich the meaning of a property in a resource. Thus, for example, the property *Area* in the concept *Parcel* is defined in &r3 as a function “*AREA[Block.BoundaryBy]*”. In this case, &r3 does not have an attribute to describe the property *Area*, but it is possible to obtain data with the same meaning applying an operator *Area[]* over *Block.BoundaryBy*. In this way, operators like *Length*, *Area*, *Buffer*, *ConvexHull*, *Union*, *Intersection*, etc. can be used. This feature can only be used when it is applied over attributes in the same resource.

With the Portal Schema and the Portal Resources Description we have all the information necessary to know what resources there are and how to satisfy a query. For example, we define the following query over the ontology shown in Figure 1: “*Obtain the Area of a Block where the Parcel(extendof) touches Block(boundary) and the Parcel(Number) > 20*”. This query is represented in our query language over a data model query language in Table 1(a).

According to Figure 1 the only resource that can satisfy all the properties defined in the query is &r3. Table 1(b) shows the results of applying translation of the properties in the ontology to real attributes in the resource &r3 (Figure 2). Note that it is a very simple translation because we use the same query language and we do not have to adapt the semantic between them. Note that the relation between *Parcel* and *Block* in the Resource &r3 is not included in the query. This relation is established on the Wrappers.

Table 1. Example of translation

SELECT Block.Area FROM Block contains Parcel WHERE Parce.Number > 20 AND Block.extendof_B TOUCHES Parcel.extendOf	SELECT AREA[B.boundaryBy] FROM Block AS B , Block.Parcel AS P WHERE P.Number>20 AND B.extendof_B TOUCHES P.extendOf
(a)	(b)

**Fig. 2.** Example of Data model – Resource &r3

Note that the queries are executed firstly over the *catalog* to filter the resources that have useful *descriptions* for the user. Finally, the user makes a spatial query over these resources. However, the description of the user interface for carrying out these queries is not the aim of this paper.

2.2 Wrapper

Generally speaking, a wrapper should execute the query and generate the results as a GML document. However, as is mentioned above, the query over GML in each source must be efficient. The spatial operator requires a spatial index that cannot be created directly over a spatial XML document. Thus, an efficient method for storing GML documents is necessary. In [13] we studied the behavior of different alternatives over XML documents (non-spatial data) applied to GML documents (spatial and non-spatial data). As a result, we selected a RDBMS model called LegoDB[14] because it obtains the best results in the study. In Figure 3, an example over the data graph in Figure 2 is shown. It is the simplest mapping proposed, usually called *inlining*.

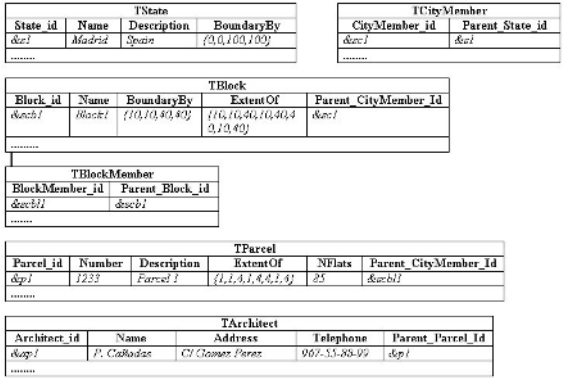


Fig. 3. Example of LegoDB approach

In view of this, we have developed a wrapper architecture to convert our query language over our data model to spatial SQL queries. This approach is not dependent on a particular DBMS, thus Oracle, Informix, DB2 or any spatial RDBMS or ORDBMS may be used. This approach can be summarized in two components: *T-Relational* and *T-DBMS*.

T-Relational transforms the query expressed in our query language (following our data model) to a query over a Relational model expressed by *LegoDB*. The output of this component is a SQL query with the syntax of our query language. This component has three important tasks: (1) conversion of attributes that contain wildcards (%,#)[12]. To do this, we need to store a structure with all the possible paths of the data model for each document that is queried. This structure is obtained from the XML Schema (DTD) of the documents; (ii) transforming each attribute of the query expressed in a data model (e.g. *block.parcel.extentof*) in the real name in the relational schema (e.g. *TParcel.extentof*); (iii) the automatic generation of *From* clause in the final SQL query, according to relations involved in the query. The relations included in the *From* clause are obtained from the attributes used in the query. In order to carry out this transformation, a definition of a mapping for the relation of our data model and the relational model is necessary. This mapping (expressed in XML) relates the semantic relation between a path in our data model with a pair (relation, attributes) in the relational model.

T-DBMS transforms the result of the *T-Relational* component to a SQL query in the syntax and semantic of a particular DBMS. This component depends on the DBMS used in the implementation. The complexity of this component depends on how to satisfy a DBMS with the OpenGIS specifications [13]. In this implementation we have used *Oracle 9i*, which follows most of the specifications defined in OpenGIS.

In Table 2 a translation of a query in Table 1(b) to a SQL query is shown. This query has the syntax of Oracle 9i. Note that the relations between relations are included.

Table 2. SQL query

```

SELECT
MDSYS.SDO_GEOM.SDO_AREA(TBLOCKS.BOUNDARY,MDSYS.SDO_DIM_ARRAY(MDSYS.SDO_DIM_ELEMENT('X', 0, 10000, 0),MDSYS.SDO_DIM_ELEMENT('Y', 0, 10000, 0)))
FROM TBLOCKS,TPARCEL,TBLOCKMEMBER
WHERE (((TBLOCKMEMBER.P_BLOCK_ID = TBLOCKS.BLOCK_ID AND TPARCEL.P_BLOCKMEMBER_ID
= TBLOCKMEMBER.BLOCKMEMBER_ID)) AND ((TPARCEL.NUMBER > 20) AND
MDSYS.SDO_RELATE(TBLOCKS.EXTENDOF,TPARCEL.EXTENDOF,'mask=TOUCH querytype =
JOIN')='TRUE'))

```

3 A Performance Study

In this section, we describe the experimental studies that were conducted in order to assess the efficiency of querying a GML document with our approach. We focus on the effectiveness of this approach in terms of translation of queries over an ontology into queries supported by the wrappers, translation of these queries into spatial SQL queries, query processing and generation of GML documents from the results. All the experiments were conducted on an 2300Mhz PC with 512Mb RAM, 40Gb hard disk with operating system Windows XP professional. The RDBMS used was *Oracle9i Spatial 9.0.1*, which we selected because it allows the storage of spatial objects and the application of spatial operators over them. Oracle's Spatial object-relational model was used (SDO_GEOMETRY object) [17] for simplicity. However, as mentioned above, Oracle's relational model or another RDBMS may be used. All experiments were conducted on the same computer, using different databases to the mediator and the wrappers in the same Oracle DBMS. In this way, we obtained a greater system load.

The data set used in these experiments respects the data set used in [13]. This data set represents a *City model* where a city has several blocks, each block has several parcels and each parcel has an owner. We used a data set with 7.1 Mb approx. and 5000 rectangular parcels. In the mediator, we used three data sets: D1 (2.1Mb approx.) with 1250 resources, D2 (4.4Mb) with 2500 resources and D3 (8.4Mb approx.) with 5000 resources.

For this test, on the DBMS the same relations shown in [13] were created. We used the simplest mapping proposed by [14]. The relations stored the data model shown in Figure 3. Indexes on relational tables were also properly built to improve query processing as follows: the attributes: *TState(state_id)*, *TCitymember(Citymember_id)*, *TBlock(block_id)*, *TBlockmember(blockmember_id)*, *TParcel(parcel_id)*, *TArquitet(arquitect_id)* are primary key. *TCitymember(parent_state_id)*, *TBlock(parent_Citymember_id)*, *TBlockmember(parent_block_id)*, *TParcel(parent_blockmember_id)*, *TArquitet(Parent_parece_id)* are foreign key. The Spatial Attributes *TState(BoundaryBy)*, *TBlock(BoundaryBy,Extendof)* and *TParcel(extendof)* have been indexed with *R-Tree* [17]. In order to carry out this mapping efficiently the RDF is stored in a RDBMS following [15].

3.1 Elapsed Time

In our study, we have focused our attention on the following aspects involved in assessing performance:

1. Translating queries over an ontology into queries supported by the wrappers has an elapsed time that should be studied [11].
2. Translating queries from our data model to the relational model (spatial SQL) needs a mapping process. It depends on the path length between the entities in the GML document [13].
3. Since we are dealing with an XML query language, a very large number of joins is necessary. It is one of the more important performance problems in the relational implementation of these query languages [18]. In addition, storing and querying spatial data requires a larger amount of resources than storing and querying alphanumeric data [19]. This elapsed time was studied in [13] in isolation from a real system, but now we study how it works inside our system.
4. Generating GML documents from the results. This process can be very expensive because a spatial query can return a very large amount of spatial objects.

For these aspects we have studied the following elapsed times: T1 is associated to the translation of queries over an ontology into queries supported by the wrappers (parser and mapping). T2 is associated to translation of these queries into spatial SQL queries (parser, mapping, generation of *from* clause and translation to SQL). T3 is the elapsed time for the execution of SQL sentences, and T4 is the elapsed time for the generation of GML documents from the result set.

Firstly, a set of queries is used to study the behavior of these approaches, involving only alphanumeric data (users may only wish to query alphanumeric data of a GML document). Secondly, a set of queries with spatial and alphanumeric operators is used. We define the complexity of a query according to the number of properties included in the query, the number of relations involved in the final spatial SQL and the number of joins in SQL. Table 4 shows the number of joins and relations involved in the final SQL query. All queries return 2500 objects approx. We have used the data set D2 in the mediator. In the execution of the query only one resource satisfies the query.

Table 3. Number of joins and relations in each query

Query	Nº Joins	Relations	Query	Nº Joins	Relations
Q1	0	1	Q7	5+1+1	6
Q2	4	5	Q8	5+1+1	6
Q3	4+1	5	Q9	3+1+1	4
Q4	5+1+1	6	Q10	3+1+1s	4
Q5	2+1	3	Q11	5+1+1s	6
Q6	5+1+1	6	Q12	5+1+1s+1	6

The elapsed query times for these queries are shown in Figure 4. Q1-Q4 are queries with only alphanumeric operators. For Q1-Q2 additional joins are not included. They have a progressive number of joins and relations. Q3-Q4 are queries with more properties in *Select* and *Where* clause and 1 additional *joins* in Q3 and 1 additional *joins* and 1 *none equi-joins* in Q4.

Q5-Q12 are queries with spatial operators, some of which have spatial and alphanumeric operators. Q5-Q6 are spatial queries with spatial operators (*Area*, *Length* and *Buffer*). The number of joins and relations is doubled in Q6.

The elapsed query times for these queries are shown in Figure 4. The elapsed time for the first mapping (mapping mediator - T1) depends on the number of resources and the number of properties involved in each query. However, the results are not very sensitive to the increment in the number of properties. The second mapping (mapping wrapper – T2) has a longer elapsed time than the first mapping. The most expensive function in this mapping is looking for all relations involved in the query and the relations between them to construct the *from clause*. However, the depth of our data model limits this elapsed time. Note that Q4 and Q6 have the highest elapsed time because they have the highest number of properties and involve entities that are highly-nested in the schema. The time of Execution (T3) of the SQL statement depends on the number of joins and the number of spatial operators involved. Note that this elapsed time increases with several joins and spatial operators. However, as is mentioned above, the number of joins can be limited by applying more specific heuristics [20]. Finally, the elapsed time to create GML documents is highest. This elapsed time increases with the adaptation of spatial objects to GML. This process was carried out following a process outside the DBMS. The process used the *Oracle sdoapi* to adapt Oracle Geometries to XML (we explore the record set to adapt the objects). In this way, we do not depend on a particular DBMS and simplify the translation to SQL. In addition, in this process a conversion between names of the relational attributes to GML attributes is carried out.

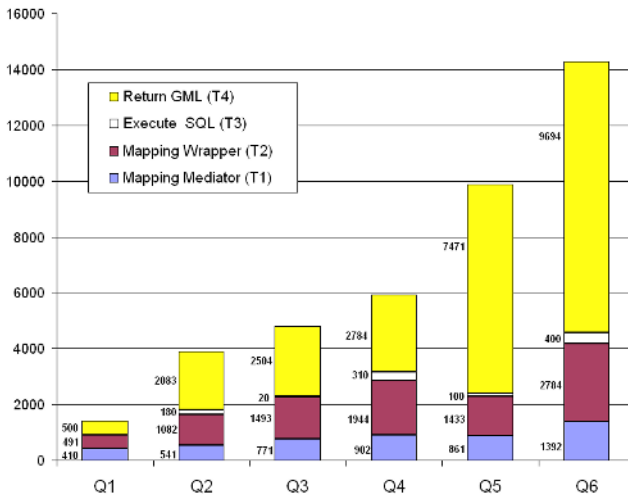


Fig. 4. Elapsed query time (Q1-Q6)

Q7-Q12 use spatial operators which are more expensive in time than Q5-Q6. Q7 uses a spatial *Union* and has a high number of joins. Q8-Q9 use spatial operators in *select* clause (*Union*, *Area*) and comparison with spatial operators (*Length*) in *where*

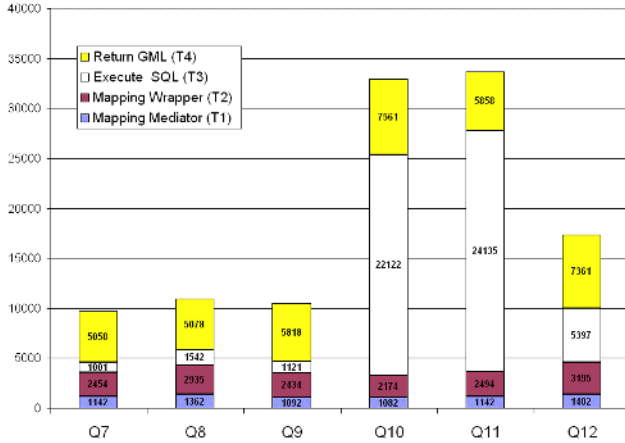


Fig. 5. Elapsed query time (Q7-Q12)

clause. Q10-Q12 are the most complex queries with a high number of joins and spatial joins. Q12 has *Union* and comparison with spatial operators (*Area*) in *where* clause for the spatial joins.

The elapsed query times are shown in Figure 5. Note that when there are complex spatial operators with a higher number of joins, execution time increases. Q7, Q8 and Q9 include, in the *select* clause, *Union* operators, combined with *Area* operator in Q8 and Q9. The elapsed time of these queries is greater than the previous set. Q10-Q12 use spatial joins (*Intersect*). Evidently, this kind of queries gets the worst elapsed time. In addition, the elapsed time to generate GML documents increases when spatial joins are involved. This is due to the process used to adapt the geometries mentioned above.

In view of this, with complex spatial queries, applying a specific study (following on from [14]) for storing GML documents over relational databases is necessary. In order to obtain a good performance querying spatial XML documents, it is preferable to increase the size of relations, and to reduce the number of joins between relations. For each source a particular heuristic must be applied, depending on the kind of information stored on them. In this way, we can reduce the number of joins in each query. This will reduce the elapsed time to execute spatial SQL and the elapsed time to get the *from* clause in the wrapper mapping. In any case, T2 plus T3 is much lower than the elapsed time for querying directly over GML [13]. However, the high elapsed time to generate GML documents could be reduced if we used specific tools offered by the DBMS (SQL/XML[17]). (This comparative is not included for reasons of space).




Fig. 6. Scalability Test

3.2 Scalability Test

With regard to scalability, three studies can be carried out: (i) scalability over the spatial Data, (ii) scalability over the mapping in the wrapper, and (iii) scalability of the translation of the query languages in the mediator. The first study does not depend on either the data in the mediator or the data in the wrapper; it just depends on the DBMS. A study with these features was carried out in [13]. The second study is not relevant because the size of the data needed to translate the queries in the translator is very small. Consequently, we only deal with the third study.

In order to do this, we used the D1, D2 and D3 data sets detailed above. Over these data sets we translated two queries with few attributes (3 attributes, Q1) and queries with a lot of attributes (15, it is not included in table 2). Thus, we can study the influence between the mapping over lots of resources involved in the query and the number of attributes of the query. In our test we suppose all resources stored in the data set satisfy the query (so, 1250 resources satisfy the queries in D1, 2500 in D2 and 5000 in D3).

Figure 6 shows the elapsed time ratios for the two kinds of queries. The elapsed query times ratios are defined as follows. Assuming the elapsed time of a query using D1 is treated as scale t_a , and the elapsed time of the same query is t_b , using either D2 or D3, the elapsed time ratio is t_b/t_a . The study shows the difference in the number of attributes is not relevant with respect to the number of resources that satisfy the query. The elapsed time of the parse is inconsiderable.

In conclusion, in our prototype the elapsed time of the mapping depends on the number of resources that satisfy the query and not on the number of attributes involved in the query. In the same way, as shown above, if there are 2500 resources but only one satisfies the query, the most significant elapsed time depends on the time to translate this query over this resource.

4 Conclusions

In this paper a prototype of a mediation system for querying XML spatial resources is studied. The main feature of this approach is the possibility of applying spatial

operators over the resources that are represented in GML format. We studied four elapsed times: regarding the elapsed time for the mapping in the mediator, we concluded that it depends on the number of resources involved in the query. For the same number of resources, the number of properties is not determinative. The elapsed time for the mapping in the wrapper is more expensive. Moreover, it is more variable because it depends on the heuristic used to represent a GML data model over the relational model in each resource. An heuristic with several relations offers a worse elapsed time in the mapping. The same problem can be applied over the execution of the query. The elapsed time for execution obtains better results with fewer relations. Finally, the worst elapsed time is obtained in the generation of GML. However, we can improve this elapsed time using specific tools offered by the DBMS.

Future work foresees the development of a prototype to allow the splitting of a query between different sources, and the joining of the result in the mediator. The study presented in this paper will be used to check the efficiency of this future approach. In addition, we are developing a prototype to integrate GML resources and any other kind of resource (HTML, PDF, etc). Thus, with this approach it is possible to discover spatial and non-spatial resources which are interrelated semantically on the Web. In this way, this work represents a small step towards the Semantic Geospatial Web.

Acknowledgment. Research Project: “Metodologías de desarrollo de interfaces de usuario dinámicas” – (JCCM PBC-03-003). Junta de Comunidades de Castilla-La Mancha.(Spain)

References

- [1] J. Egenhofer. Toward the Semantic Geospatial Web. ACM-GIS 2002. 10th ACM International Symposium on Advances in Geographic Information Systems. McLean (USA). 2002.
- [2] OpenGis Consortium. Specifications. <http://www.opengis.org/techno/specs.htm>.1999
- [3] H. Samet. Applications of spatial data structures. Computer Graphics, Image processing and GIS. Addison – Wesley. 1990.
- [4] J. Córcoles and P. González. Querying Spatial Resources. An Approach to the Semantic Geospatial Web. CAiSE'03 workshop "Web Services, e-Business, and the Semantic Web (WES): Foundations, Models, Architecture, Engineering and Applications". To Appear in Lecture Notes in Computer Science (LNCS) by Springer-Verlag. 2003.
- [5] OpenGIS. Geography Markup Language (GML) v3.0. <http://www.opengis.org/techno/documents/02-023r4.pdf>. 2003
- [6] B. Amann, I. Fundulaki and M. Scholl, C. Beeri, A-M. Vercoustre. Mapping XML Fragments to Community Web Ontologies. In Proc. Fourth International Workshop on the Web and Databases. 2001
- [7] B. Amann, C. Beeri, I. Fundulaki, and M. Scholl. Ontology-Based Integration of XML Web Resources. In *International Semantic Web Conference (ISWC)*, Sardinia, Italy, 2002.
- [8] O. Boucelma, M. Essid and Z. Lacroix. A WFS-Based Mediation System for GIS Interoperability. ACM-GIS 2002. 10th ACM International Symposium on Advances in Geographic Information Systems. McLean (USA). 2002

- [9] A. Gupta, R. Marciano, I. Zaslavsky and C. Baru. Integrating GIS and Imagery through XML based information Mediation. *Integrated Spatial Databases: DigitalImages and GIS. Lecture Notes in Computer Science. Vol1737. Pp. 211-234. Springer-Verlag. 1999*
- [10] J. Córcoles, P. González and V. López-Jaquero. Integration of Spatial XML Documents with RDF. *International Conference on Web Engineering. Spain. To Appear in Lecture Notes in Computer Science (LNCS) by Springer-Verlag. 2003.*
- [11] A.Y. Levy. Logic-Based Techniques in Data Integration. In Jack Minker, editor, *Logic Based Artificial Intelligence*, pages 575-595. Kluwer, 2000.
- [12] J. Córcoles and P. González. A Specification of a Spatial Query Language over GML. *ACM-GIS 2001. 9th ACM International Symposium on Advances in Geographic Information Systems. Atlanta (USA). 2001*
- [13] J. Córcoles and P. González. Analysis of Different Approaches for Storing GML Documents *ACM-GIS 2002. 10th ACM International Symposium on Advances in Geographic Information Systems. McLean. (USA) 2002*
- [14] P. Bohannon, J. Freire, P. Roy and J. Simeon. From XML Schema to Relations: A Cost-Based Approach to XML Storage. *18th International Conference on Data Engineering (ICDE2002). 2002.*
- [15] S. Alexaki, V. Christophides, G. Karvounarakis, D. Plexousakis, K. Tolle. “*The ICSFORTH RDFSuite: Managing Voluminous RDF Description Bases*”. In *Proceedings of the 2nd International Workshop on the Semantic Web (SemWeb’01)*, in conjunction with WWW10, pp. 1-13, Hong Kong., 2001.
- [16] Karvounarakis, G., Christophides, V., Plexousakis, D., and Alexaki, S. Querying community web portals. Technical report, Institute of Computer Science, FORTH, Heraklion, Greece. 2000 <http://www.ics.forth.gr/proj/isst/RDF/RQL/rql.pdf>.
- [17] Oracle9i Database Documentation. <http://otn.oracle.com:80/docs/products/oracle9i/content.html>. 2002.
- [18] S. Abiteboul, P. Buneman and D. Suciu. *Data on the Web. From Relations to Semistructured Data and XML.* Morgan Kaufmann Publishers. 2000.
- [19] P. Rigaux, M. Scholl and A. Voisard. *Spatial Databases with Application to GIS.* Morgan Kaufmann Publishers. 2002.
- [20] J. Shanmugasundaram, K. Tufte, G. He, C. Zhang, D. DeWitt, and J. Naughton. Relational Databases for Querying XML Documents: Limitations and Opportunities. In *Proc. of the Int’l. Conf. On Very Large Data Bases*, pages 3002-314. 1999.