3rd International IEEE Conference Intelligent Systems, September 2006

Towards the Practical Use of Qualitative Spatial Reasoning in Geographic Information Retrieval

Alia I. Abdelmoty, Philip Smart, Baher A. El-Geresy

Abstract— **Geo-ontologies have a key role to play in the development of the geo-semantic web, with regard to facilitating the search for geographical information and resources. They normally hold large amounts of geographic information and undergo a continuous process of revision and update to ensure their currency. Hence, means of ensuring their integrity are crucial and needed to allow them to serve their purpose. This paper proposes the use of qualitative spatial reasoning as a tool to support the development of a geo-ontology management system. Spatial integrity rules based on uniform and hybrid spatial reasoning are proposed for the automatic derivation of spatial relationships and for maintaining the spatial consistency of the geographic data. A framework for the representation of and reasoning over geo-ontologies is presented using the web ontology language OWL and its associated reasoning tools. Spatial reasoning and integrity rules are represented using a spatial rule engine extension to the reasoning tools associated with OWL. To demonstrate the proposed approach, a case study showing a prototype geo-ontology and the implementation of the spatial reasoning engine is presented. This work is a step towards the realisation of a complete geo-ontology management system for the semantic web.**

Index Terms— **Spatial Reasoning, Geographic Information Retrieval, Geo-ontologies, Geo-semantic web.**

I. INTRODUCTION

Retrieval of geographically-referenced information on the internet is now a common activity. A large number of documents stored and retrieved on the web include references to geographic information, typically, by means of place names. Also, the web is increasingly being seen as a medium for the storage and exchange of geographic data sets in the form of maps. The geo-semantic web (GeoWeb) is being developed to address the need for access to current and accurate geo-information [Ege02]. The potential applications of the GeoWeb are numerous, ranging from specialized application domains for storing and analyzing geo-information to more common applications by casual users for querying and visualizing geo-data, e.g. finding location of services, descriptions of routes, etc.

At the heart of the GeoWeb are geographic ontologies or geo-ontologies. These are models of terminology and structure of geographic space as well as records of entities in this space. An example of such an ontology has been proposed recently in the SPIRIT project [JAF$^+$04], [FJA05] and was shown to play a central role in the development of a geographical search engine. Building geo-ontologies involves a continuous process of update to the originally modelled data to reflect change over time as well as to allow for ontology expansion by integrating new data sets, possibly from different sources. One of the main challenges in this process is finding means of ensuring the integrity of the geo-ontology and maintaining its consistency upon further evolution. Developing methods for the management of the spatial integrity of geo-ontologies will contribute towards the development of reliable geographical search engines and to the success of the GeoWeb in general.

In this paper we propose a new framework for the management of geo-ontologies for the purpose of geo-information retrieval. In particular, we build upon and utilise research results in the area of qualitative spatial reasoning (QSR). Composition tables for different types of qualitative spatial relations are used to derive general rules that govern the structure of the geographic entities and their interaction in space. A spatial integrity rule language has been developed, as an extension to the web ontology language OWL, for the expression of these rules. OWL and its associated reasoning engine Jena are used for the representation and reasoning over the geo-ontology. A demonstration of the application of the proposed framework is given using a sample geo-ontology.

A. *Integrity in Space*

Inaccuracy or error in geographic data can be accumulated at different stages of handling and using the data [ea99], from the data collection phase, to maintenance and update processes on stored data. Errors in the description of the location and shape of geo-objects can propagate to errors in the spatial relationships between those objects, and consequently to wrong information being retrieved and analysed by users. For example, error in an object's location may lead to it crossing another object when it should have been adjacent to it. Spatial relationships between geo-objects, recognized by visual interpretation, are mostly implicit and are only derived when needed using geometric computations. Erroneous updates to geographic data sets may go undetected unless appropriate spatial integrity rules are declared and applied.

Taxonomies of spatial integrity constraints have been proposed in [ea00c], [Coc04], [AJ97]. Geometrical, topological, topo-semantic and user-defined constraints have been identified. Besides basic geometric validation in the process of data cleaning carried out in some GIS, few works address the formal treatment of spatial integrity constraints [ea00a], [ea00b] and no methodology has yet been adopted in conventional GIS or spatial database management systems. Scenarios for some types of topological error correction are proposed in [UE97].

Alia I. Abdelmoty and Philip Smart are with the Cardiff School of Computer Science, Cardiff University, Wales, UK (e-mail: a.i.abdelmoty,psmart@cs.cf.ac.uk)

Baher A. El-Geresy is with the School of Computing, University of Glamorgan, Wales, UK (e-mail: bageresy@glam.ac.uk)

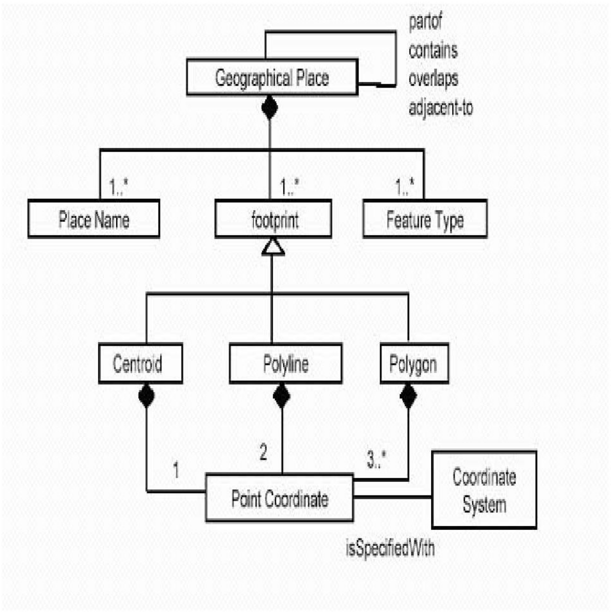

Fig. 1. Conceptual design of the geo-ontology [JAF+04].

B. *Geo-ontologies*

In this paper we are concerned with supporting geo-ontologies that are used to serve the purposes of geographical information retrieval on the web. An example of such an ontology was proposed recently by the SPIRIT project [JAF+04]; a project concerned with the development of a geographically-aware web search engine. The geo-ontology in SPIRIT provided a model of the terminology and structure of the geographic space and played a key role in supporting the various components of the system, including, interpretation of user queries in the user interface, the formulation of system queries, generation spatial indexes to support spatial search, relevance ranking of query results as well as geographic metadata extraction from web resources [FJA05]. The main concept in this ontology, as shown in figure 1 is a geographic place that is associated with one or more place names. Multiple types of spatial relations are supported by the ontology, including containment, adjacency and directional relationships.

In section 2, a new framework is proposed for the representation and management of the geo-ontology. A typology of spatial integrity constraints is presented in section 4 with a discussion of how uniform and hybrid QSR can be used for the identification of reasoning rules over geo-ontologies. An overview of the implementation of the spatial integrity rules ausing OWL and Jena is illustrated with a case study in section 5.

II. Ontology Maintenance System Architecture

The current version of the SPIRIT system uses the Oracle Spatial database system for storing the geo-ontology. Oracle Spatial provides support for spatial data types and geometric operations for the processing of functions such as intersection of polygons and distance calculations, etc. The extensibility of this system is limited for the following reasons.

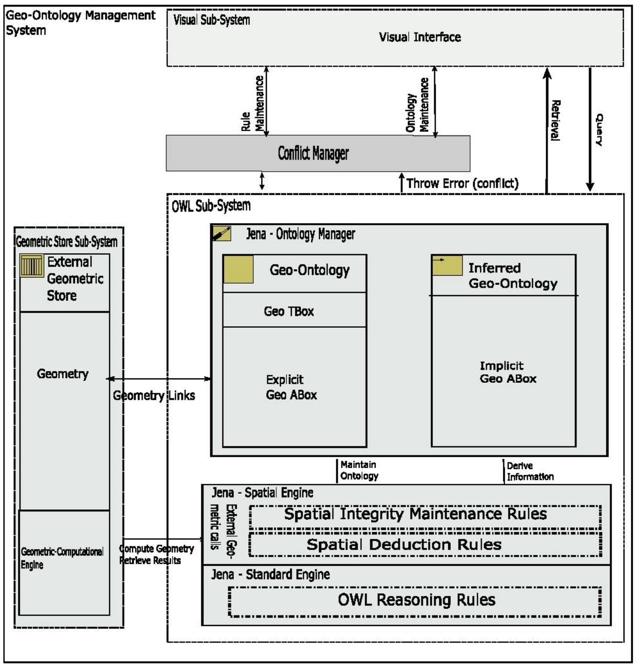

Fig. 2. New framework for representing and maintaining the geo-ontology.

- The correctness of the ontology and its viability are dependent on the quality of the data input procedure, and will suffer if appropriate validation checks of spatial consistency are missed out.
- Any update procedures on the ontology, e.g. integrating new data sets, may result in violating the consistency of the ontology and consequently degrading its fitness for use.

Here we propose an alternative framework for supporting the geo-ontology and enhancing the extensibility of systems employing such ontologies. The new framework is shown in figure 2.

The new framework utilises the web ontology language OWL and the Jena toolkit for storing and maintaining the geo-ontology. Follows is a description of its components.

Owl geo-ontology store: Both the geo-ontology model as well as its realised instances are represented and stored in OWL (Geo-TBOX and Geo-ABOX respectively). The geometric component of the geographic features, representing actual coordinates, are normally very large and constitutes a significant proportion of the size of the data files. Current OWL reasoners are not capable of handling large data sets and hence, these geometric components are stored separately in another system (e.g. Oracle Spatial). Appropriate links are made between the OWL ontology and the geometric component to facilitate retrieval and update processes. This framework also allows for appropriate delegation of the processing of geometric computational procedures, e.g. area and distance calculations to the geometric data store. Such computations are more effectively handled in a specialised system outside Owl.

Spatial Rule Engine: The rule engine is used to represent and process different types of spatial rules, namely, deductive rules and integrity rules. Deductive rules are used to automatically extract new, implicit, spatial properties and relationships from the geo-ontology. The inferred facts will be stored explicitly, in the inferred geo-ontology store (implicit Geo-ABox), and will be used to facilitate the querying and maintenance of the existing store. Spatial integrity rules are also supported and used for checking the consistency of the existing ontology and identifying any inconsistencies on update. This component interacts with an external geometric data processor for evaluating any geometric computations required. Spatial rules are discussed in detail later.

Standard Rule Engine: This is the standard Jena engine used for reasoning with the underlying semantics of OWL. It will be used to derive the entailments of the OWL language (subsumption and property reasoning). Derived facts will also be added to the inferred ontology store.

Conflicts Manager: This component is used to identify and filer out conflicts on updates to the ontology and the rule engine. The source of conflict including the objects and relationships involved as well as the fired integrity rules are recorded and reported to the user interface. Also, possible alternative update scenarios may be computed and suggested and through interaction with the user more suitable options are chosen and propagated back to the ontology.

III. A Typology of Spatial Integrity Rules

Three types of spatial integrity rules have been classified previously in the literature, namely, topological, topo-semantic and user-defined rules. Topological relationships is only one type of possible other types of spatial relationships that may be employed in devising integrity constraints for spatial data. Here we identify basic space rules derived by observing properties and relationships of objects in space.

A. Basic Laws of Space

The structure of space formed by the locations, properties and relationships between the objects that exist in it are governed by "commonsense" laws that determine its integrity and suggest its feasibility. Many of these "space laws" are implicitly recognised and learnt by humans and give us our ability to determine and infer the logical structure of space. For example, the fact the one object contains another implies that it is also larger in size. Other laws are not as obvious, but are derived using mathematical and geometrical techniques, for example, the fact that a square whose area is equal to that of a circle will not fit inside the circle.

In this work we propose to derive and make use of those space laws and translate them explicitly into integrity rules that can be used to maintain the consistency of spatial data sets. A simple classification of space laws can be between those derived from object properties and those based on spatial relationships. In what follows, an overview is presented of the types of rules from both categories.

A.1 Space Laws Based on Object Properties

Inherent properties of spatial objects derived from their shape and size can be expressed as constraints which those objects should conform to in space. Some examples of this type of rules for simple geometric data types include the following.

- A region, represented by a simple polygon, must have at least three different points.
- A simple polygon must be closed, i.e. its last point is the same as the first point.
- A simple line has at least two different points.

Other examples for more complex spatial data types include the following.

- A network is formed of a set of connected line segments.
- The curvature of line segments in a road network is normally in a specific range.

The above are examples of possible constraints that may be used to force the integrity of the objects' shapes. In the following, constraints derived from the interaction of objects in space are studied.

A.2 Space Laws Based on Spatial Relationships

A large body of research has been undertaken in the past decade in the field of qualitative spatial reasoning (QSR) with the main aim of deriving compositions of spatial relations between different objects in space. Several works studied the representation of different types of spatial relationships and complete and sound sets of relationships have been reported for different types of simple geometric shapes. In what follows, uniform reasoning with single types of spatial relationships are first presented, followed by hybrid reasoning where different types of relationships are used to encode spatial reasoning rules.

A.3 Uniform reasoning

Generalised composition table for simple regions and regions and lines are shown in tables I and II respectively. Composition tables are used to record the results of the interaction of spatial relations between objects. The tables record the relationships between three spatial entities, A, B, and C. Entries in the table are sets of relationships that may result from the composition of the corresponding raw and column entries. For example, given the fact that $inside(A, B)$ and $inside(B, C)$, the fact that $inside(A, C)$ can be inferred. Recently, methods were proposed for the homogeneous representation of different types of spatial relationships [EGA02]. Also, associated reasoning mechanisms for the composition of spatial relations and the automatic production of composition tables have been presented in [EGA04].

Entries in the composition tables can be encoded into rules with a view of playing two roles; firstly as deduction rules for the automatic derivation of implicit spatial relationships, and secondly as constraints for enforcing the integrity of the spatial data sets.

Rules for reasoning over orientation as well as proximal relationships have also been studied. Table III shows the

	$d(y, z)$	$m(y, z)$	$i(y, z)$	$ct(y, z)$	$o(y, z)$
$d(x, y)$ 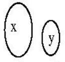	all	$d \vee m \vee$ $i \vee o$	$d \vee m \vee$ $i \vee o$	d	$d \vee m \vee$ $i \vee o$
$m(x, y)$ 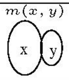	$d \vee m \vee$ $ct \vee o$	$d \vee m \vee$ $i \vee ct \vee o$	$i \vee o$ o	d	$d \vee m \vee$ $i \vee o$
$i(x, y)$ 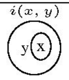	d	d	i	all	$d \vee m \vee$ $i \vee o$
$ct(x, y)$ 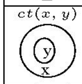	$d \vee m \vee$ $ct \vee o$	$ct \vee o$	$i \vee ct \vee$ o	ct	$ct \vee o$
$o(x, y)$ 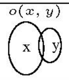	$d \vee m \vee$ $ct \vee o$	$d \vee m \vee$ $ct \vee o$	$i \vee o$	$d \vee m \vee$ $ct \vee o$	all

TABLE I

Composition table for the set of base topological relations between simple regions.

	$d(x, y)$ 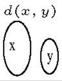	$m(x, y)$ 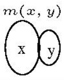	$i(x, y)$ 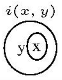	$ct(x, y)$ 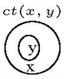	$o(x, y)$ 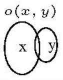
$disjoint(y, z)$ 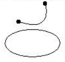	all	all	d	all	all
$touch(y, z)$ 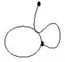	$d \vee m \vee o$	all	$d \vee m$	$i \vee o$	all
$inside(y, z)$ 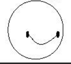	d	$d \vee m$	all	i	all
$overlap(y, z)$ 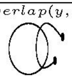	$d \vee m \vee o$	$d \vee m \vee o$	$d \vee m \vee o$	$i \vee o$	all

TABLE II

The composition table between two regions and a region and a line.

composition table for the cardinal direction relationships between two simple regions.

A.4 Hybrid Spatial Reasoning

As stated earlier, four different types of distinct qualitative spatial relationships are identified. The composition tables presented so far consider individual relationships in isolation of others. However, the interaction of objects in space can in fact be expressed as a mixture of different relationships. For example, one object is overlapping and smaller than another, or an object is adjacent to and north of another, etc. Building composition tables using more than one spatial relationship should provide more richer and definite rules. Table IV shows the modified topological composition table between two regions given knowledge of the relative size relationships between the objects. As can be seen in the figure, some entries are now invalid (shown empty in the table), and other entries are refined.

	N	NE	E	SE	S	SW	W	NW
N	N	N	NE				NW	NW
NE	NE	NE	NE					
E	NE	NE	E	SE	SE			
SE			SE	SE	SE			
S			SE	SE	S	SW	SW	
SW					SW	SW	SW	
W	NW				SW	SW	W	NW
NW	NW						NW	NW

TABLE III

(B) Definite results in their composition table. O_c represents a neutral zone (Adapted from [Frank 92]).

$x > z \wedge$	$d(y, z)$	$m(y, z)$	$i(y, z)$	$ct(y, z)$	$o(y, z)$
$d(x, y)$	all	$d \vee m \vee$ o	$d \vee m \vee$ o	d	$d \vee m \vee$ o
$m(x, y)$	$d \vee m \vee$ $ct \vee o$	$d \vee m \vee$ $ct \vee o$	o	d	$d \vee m \vee$ o
$i(x, y)$	d	d		all	$d \vee m \vee$ $i \vee o$
$ct(x, y)$	$d \vee m \vee$ $ct \vee o$	$ct \vee o$	$ct \vee o$	ct	$ct \vee o$
$o(x, y)$	$d \vee m \vee$ $ct \vee o$	$d \vee m \vee$ $ct \vee o$	o	$d \vee m \vee$ $ct \vee o$	all

TABLE IV

Hybrid topological-size composition table. The table shows the resulting compositions with the knowledge that $x > y$.

IV. Overview of the Spatial Rule Engine

Jena2 is an open source Java-based semantic web tool. It provides an API to access, manipulate and reason with RDF, DAML+OIL and OWL ontologies. Jena can be used to reason with OWL using OWL's standard inference mechanisms and to construct custom-made rules. An extension to the OWL rule engine in the Jena toolkit is developed here to realise the spatial reasoning rules.

The main functions of the extended rule engine is to allow for the execution of integrity rules over the geo-ontology to check for inconsistencies, as well as to allow for the deduction of implicit information and the building of the inferred Geo-ABox as shown in figure 2.

Jena uses the RETE pattern matching algorithm [Sch92] for matching variables of atoms. However, the rule engine will need to rely on external geometric processors, such as that in the Oracle Spatial database system for the computation and evaluation of geometric facts. Hence, some modification of the system is needed to allow for such external calls to be integrated. Currently result of calls take the form of either a boolean value, a matched individual (RDF:ID) or an numerical value. Also, plans for the scheduling and execution of rules and for consequent firing need to be devised to determine the order of rule application in case of conflicts [AW03]. The rule engine will interact with the user interface through the conflict manager for error reporting.

To demonstrate the functionality of the rule engine, a case study is given below of a sample geo-ontology. The representation of the ontology and examples of spatial integrity maintenance are presented. A detailed description

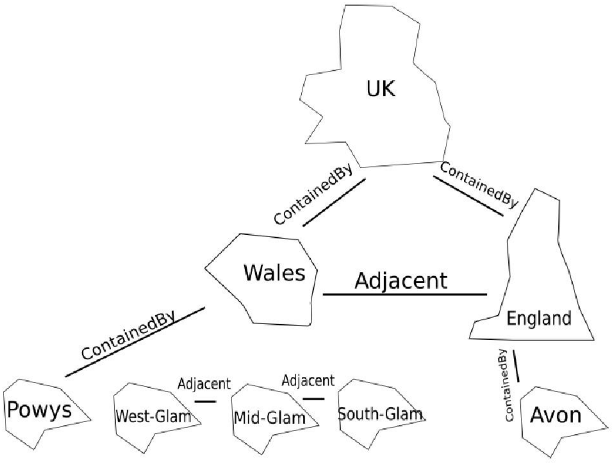

Fig. 3. Use Case geographic Scene

of the rule engine and the underlying rule language is beyond the scope of this paper.

A. *Case Study*

A sample geo-Ontology with a simple model is created using the OwlViz [Hor] plugin to Protege [GEF$^+$99], where the only class is a region, and the only properties the different types of relationships between regions. The ontology defined for this case study is shown in figure 3.

A sample of the rules implemented in the engine in XML syntax is as follows.

```
[RuleID = "EQ-Inside"; RuleType="Topological": Region(?A)
Region(?B) Region(?C) EQ(?a ?b)Inside(?b ?c)
 isValid(Inside ?a ?c) -> throw(?a ?b ?c) ]

[RuleID = "Inside-Disjoint" ;RuleType="Topological": Region(?A)
Region(?B) Region(?C) Inside(?a ?b) Disjoint(?b ?c)
isValid(Disjoint ?a ?c) -> throw(?a ?b ?c) ]

[RuleID = "TRANS-Inside"; RuleType="Topological": Region(?A)
Region(?B) Region(?C) Inside(?a ?b) Inside(?b ?c)
 isValid(Inside ?a ?c) -> throw(?a ?b ?c) ]

[RuleID = "SYM-Adjacent": RuleType="Topological": Region(?A)
Region(?B) Adjacent(?A ?B)
 isValid(Adjacent ?B ?C) -> throw(?A ?B ?C) ]
```

The isValid(rel ind ind) spatial function is used to check the consistency of a relationship between the individuals bound to ?A and ?C. The engine uses a constraint satisfaction algorithm to perform this check. If an error is found, the individuals bound in the rule will be thrown. Three possible errors may occur.

- The original relationship between the individuals bound to ?A and ?C is invalid.
- The relationship between the individuals bound to ?A and ?B is invalid.
- The relationship between the individuals bound to ?B and ?C is invalid.

A visual interface is built to visualize the outcome of the execution of the rule language through the prototype engine. The tool visualizes the regions in the underlying OWL ontology model and the relationships between them. Regions are denoted by circles and relationships are denoted by lines. If an integrity rule is violated, the concerned relationship(s) edge(s) as well as the objects involved are highlighted. In the reminder of this section an update scenario to the sample ontology above will be demonstrated.

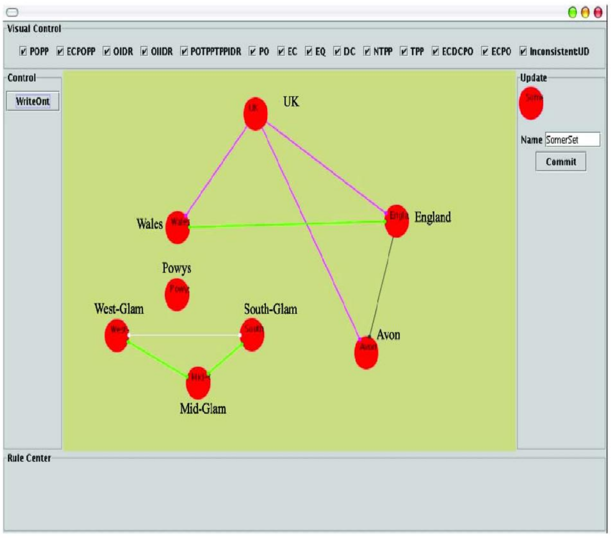

Name SomerSet

Fig. 4. Visual interface showing the initial state of the geo-ontology

The Initial Sample Scene

The initial scene is depicted in figure 3. The ontology is populated with the following individuals (regions); Wales, UK, England, Avon, South-Glam, Mid-Glam, West-Glam, Powys. The following relationships are explicitly stored.

```
 Wales adjacent England. Wales inside UK.
England inside UK. South-Glam adjacent Mid-Glam.
 Mid-Glam adjacent West-Glam. Avon inside England.
```

This initial state is shown by the visual interface in figure 4. The following are XML/RDF OWL fragments of the full ontology which has been constructed with respect to the scene shown in figure ??.

```
<rdf:Description rdf:about="http://www.geo.ont/#Region">
  <rdf:type rdf:resource="http://www.w3.org/2002/07/owl#Class"/>
</rdf:Description>
```

The properties used in the ontology are shown below

```
<rdf:Description rdf:about="http://www.geo.ont/#Disjoint">
 <rdfs:range rdf:resource="http://www.geo.ont/#Region"/>
 <rdfs:domain rdf:resource="http://www.geo.ont/#Region"/>
 <rdf:type rdf:resource="http://www.w3.org/2002/07/
   owl#ObjectProperty"/>
</rdf:Description>

<rdf:Description rdf:about="http://www.geo.ont/#Adjacent">
 <rdfs:range rdf:resource="http://www.geo.ont/#Region"/>
 <rdfs:domain rdf:resource="http://www.geo.ont/#Region"/>
 <rdf:type rdf:resource="http://www.w3.org/2002/07/
   owl#ObjectProperty"/>
 </rdf:Description>
```

The individuals of type Region namely; Wales, the UK and Avon are shown below:

```
<rdf:Description rdf:about="http://www.geo.ont/#Wales">
  <j.0:Adjacent rdf:resource="http://www.geo.ont/#England"/>
  <rdf:type rdf:resource="http://www.geo.ont/#Region"/>
  <j.0:Inside rdf:resource="http://www.geo.ont/#UK"/>
```

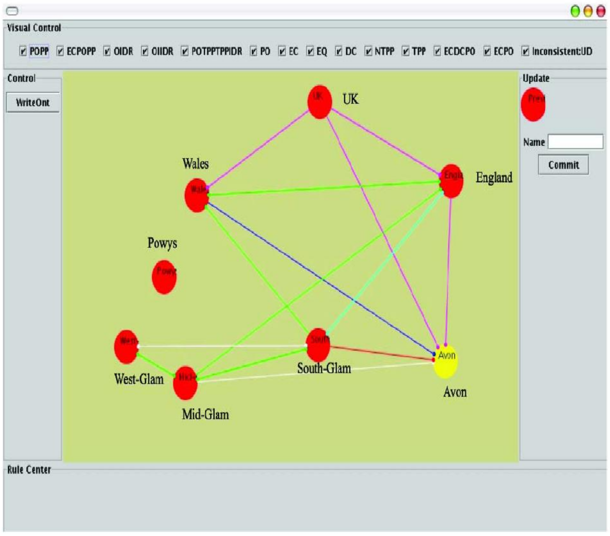

Fig. 5. Integrity violation in the update scenario.

```
</rdf:Description>

<rdf:Description rdf:about="http://www.geo.ont/#UK">
  <rdf:type rdf:resource="http://www.geo.ont/#Region"/>
</rdf:Description>

<rdf:Description rdf:about="http://www.geo.ont/#Avon">
  <j.0:Inside rdf:resource="http://www.geo.ont/#England"/>
  <rdf:type rdf:resource="http://www.geo.ont/#Region"/>
</rdf:Description>
```

Update Scenario

Step 1 Update the ontology by adding the fact that Avon is disjoint from South-Glam.

Step 2 Adding the fact that England is equal to South-Glam causes an integrity violation as shown in figure 5.

Violation The violation detected is between Avon and South-Glamorgan (Avon is inside England and if England is equal to South Glamorgan, then Avon should also be inside South Glamorgan). This is indicated on the visual interface as a highlighted line linking the two regions. The integrity rule that has been violated is the "EQ-Inside" rule defined above.

V. CONCLUSIONS

This paper is concerned with the issue of development and management of geo-ontologies on the semantic web. A framework for a geo-ontology maintenance system was proposed that uses the web ontology language OWL for the representation of the geo-ontology and extends its reasoning engine with a spatial rule engine for expressing and implementing spatial integrity maintenance rules. Qualitative spatial reasoning using one type of spatial relationships as well as hybrid reasoning using multiple types of relationships was used to identify spatial constraints in the system. The approach proposed is a step towards the development of a complete geo-ontology management system for the Geo-semantic web. Further work is in progress to evaluate the validity of the approach proposed using realistic data sets in large geo-ontologies.

REFERENCES

[AJ97] A.I. Abdelmoty and C.B. Jones. Towards Maintaining Consistency of Spatial Databases. In *CIKM'97*, pages 293–300. ACM press, 1997.

[AW03] Grigoris Antoniou and Gerd Wagner. Rules and defeasible reasoning on the semantic web. In *RuleML*, pages 111–120, 2003.

[Coc04] Sophie Cockcroft. The design and implementation of a repository for the management of spatial data integrity constraints. *Geoinformatica*, 8(1):49–69, 2004.

[ea99] P.A. Longley et al. *Geographical Information Systems.* John Wiley and Sons, 1999.

[ea00a] A. Belussi et al. An integrity constraints driven system for updating spatial databases. In *8th ACM GIS*, pages 121 – 128. ACM Press, 2000.

[ea00b] K.A.V. Borges et al. Integrity Constraints in Spatial Databases. In *Database integrity: challenges and solutions*, pages 144 – 171. ACM Press, 2000.

[ea00c] S. Servigne et al. A Methodology for Spatial Consistency Improvement of Geographic Databases. *GeoInformatica*, 4(1):7–34, 2000.

[EGA02] B.A. El-Geresy and A.I. Abdelmoty. Towards a General Theory for Modelling Qualitative Space. *International Journal on Artificial Intelligence Tools, IJAIT*, 11(3):347–367, 2002.

[EGA04] B.A. El-Geresy and A.I. Abdelmoty. SPARQS: a Qualitative Spatial Reasoning Engine. *Journal of knowledge-based Systems*, 17(2-4):89–102, 2004.

[Ege02] Max J. Egenhofer. Toward the semantic geospatial web. In *Proceedings of the tenth ACM international symposium on Advances in geographic information systems*, pages 1–4. ACM Press, 2002.

[FJA05] G. Fu, C.B. Jones, and A.I. Abdelmoty. Ontology-based spatial query expansion in information retrieval. In *On the Move to Meaningful Internet Systems 2005: ODBASE 2005*, volume LNCS 3761, pages 1466 – 1482. Springer Verlag, 2005.

[GEF$^+$99] W. Grosso, H. Eriksson, R. Fergerson, J. Gennari, S. Tu, and M. Musen. Knowledge modeling at the millennium (the design and evolution of protege-2000). In *Proceedings of the 12 th International Workshop on Knowledge Acquisition, Modeling and Mangement (KAW'99)*, 1999.

[Hor] Matthew Horridge. Owlviz. http://www.co-ode.org/downloads/owlviz/co-ode-index.php.

[JAF$^+$04] C.B. Jones, A.I. Abdelmoty, D. Finch, G. Fu, and S. Vaid. The spirit spatial search engine: Architecture, ontologies and spatial indexing. In *Geographic Information Science: 3rd Int. Conf., (GIScience04)*, volume LNCS 3234, pages 125–139. Springer Verlag, 2004.

[Sch92] Bruce Schneier. The rete matching algorithm. *AI Expert*, 7(12):24–29, December 1992.

[UE97] T. Ubeda and M.J. Egenhofer. Correcting Topological Errors. In *5th SSD'97*, volume 1262, pages 283–297, 1997.