Spatial integration of Semantic Web Services: the e-Merges approach

Conference or Workshop Item

Spatial Integration of Semantic Web Services: the e-Merges Approach*

Vlad Tanasescu[1], Alessio Gugliotta[1], John Domingue[1],
Leticia Gutiérrez Villarías[2], Rob Davies[2], Mary Rowlatt[2], Marc Richardson[3],
Sandra Stinčić[3]

[1] Knowledge Media Institute, The Open University,
Walton Hall, Milton Keynes, MK7 6AA, UK
{v.tanasescu, a.gugliotta, j.b.domingue}@open.ac.uk

[2] Essex County Council, County Hall,
Chelmsford, CM1 1LX, UK
{Leticia.gutierrez, maryr}@essexcc.gov.uk
rob.davies@mdrpartners.com

[3] BT Group
Adastral Park Martlesham, Ipswich IP5 3RE, UK
{marc.richardson, sandra.stincic}@bt.com

Abstract. As Semantic Web Services (SWS) are becoming a more mature technology, the question of their integration into the web landscape is pushed to the foreground. In a world where it is believed that up to 80% of data has a geographical component, one in which new web maps applications recently show tremendous growth, and in which of course we constantly think and act in terms of movement and geographic features, integration into the spatial domain appears as an essential step toward wide-scale adoption of SWS technology. However, geographic space, as a unique but all encompassing domain has specificities that semantic descriptions must acknowledge. Furthermore, Geographical Information Systems (GIS) need to adapt to human cognitive abilities of spatial representation and reasoning. In this context, e-Merges, an emergency management application prototype developed in collaboration with emergency planners of public agencies, is an ongoing effort to integrate SWS technology in a GIS environment, by applying the SWS notions of goal and context based interaction.

Introduction

Semantic Web Services (SWS) are the result of an acknowledgement that Web Service technology (WS), even in its standardized form, cannot achieve a satisfying

* This work was supported by the DIP (Data, Information and Process Integration with Semantic Web Services) project. DIP (FP6 - 507483) is an Integrated Project funded under the European Unions IST programme.

level of interoperability without appropriate high-level semantics. Indeed, WS based on standards such as UDDI[1] for discovery, WSDL[2] for interface description, and SOAP[3] for message passing, simplify the task of the developer but without dismissing his or her knowledgeable intervention. Indeed, when new services are to be integrated to an application, developers need to study the WS descriptions, to match inputs, outputs and invocation workflows with the existing systems.

On the other hand, by using SWS, if the vision of fully automatic interaction and composition is still somehow remote, the following tasks are greatly alleviated:

- *Discovery* of useful services is achieved by matching a formal task description against SWS' semantic descriptions.
- *Mediation* between heterogeneous services can be specified at the level of data format, message protocol and business processes.
- *Composition* of services provides a means of creating a new service by aggregating existing components.

IRS-III [1], a platform and broker for developing and executing semantic Web services, adopts a semantic Web approach based on ontological descriptions, expressed formally in OCML [7]. In particular IRS-III incorporates and extends the Web Services Modeling Ontology (WSMO) [8]. *Goals*, a concept existing in WSMO, can be invoked in this extension, which ensures an intuitive way of interacting with clients in a Semantic Web (SW) context.

In parallel to the SW and SWS efforts, another, maybe more spontaneous, evolution of the web is taking place. Web2.0 applications, by offering large amounts of resources to users for small fees, weaving large social networks where previously only forests of text based hyperlinks existed, and providing desktop like applications in the browser, are changing the way we interact on the Web. Part of this evolution is a renewal of the available mapping applications; closed, static and schematic traditional web map applications are progressively replaced by *new web maps*, intensively using AJAX technologies and employing new means to achieve what we call the *map reality effect*, an effort of rooting the maps into the cognitive reality by giving more natural looking insights into the geography covered by it. Also, by freely distributing APIs, new web maps lead to an explosion of *mashups*, minimal applications developed by independent technically skilled users which aggregate data on a spatial context in order to fulfill a specific goal.

This last evolution more than everything else shows the interest and the appeal of the spatial context for web users; mashups are used for a wide variety of goals, to such extent that it seems that space, mediated through realistic web maps, came to represent a link between the mostly textual world of the internet and the World itself. Indeed, geographic space may provide the terrain for data integration rooted into human cognition that the more abstract textual web seems to fail to achieve. It also provides a link with the user's daily existence and situated context, which may allow intelligent filtering of otherwise overwhelming resources, while not restraining their accessibility.

[1] http://www.uddi.org/
[2] http://www.w3.org/TR/wsdl
[3] http://www.w3.org/TR/soap12-part1/

To acknowledge these evolutions, we applied SWS technologies to the spatial domain in the e-Merges prototype, which has been designed as an e-Government use case in the context of the DIP project (funded under the European Union's IST programme FP6). E-Merges illustrates the way in which spatially related data (SRD) delivered through SWS can ease the management of a specific use case by aggregating data originating from different sources, and presenting it in a way which is consistent and task relevant.

We first present our generic approach to the representation of objects and context in the spatial domain, then explain how SWS applications are build using IRS-III, and finally, before concluding, describe aspects of the e-Merges prototype.

Semantics for the Geographic Space

It is well acknowledged that the spatial domain is somehow *special* [9]. Indeed, Geographic Space encompasses objects quite different from the ones we usually manipulate or are used to describing in knowledge bases; scale, orientation, boundaries, and cultural conceptions, amongst other elements, seem to matter to a greater extent [10].

Therefore, if a full review of the specificity of the geographic domain is beyond the scope of our work, three aspects of this specificity particularly oriented our research:

- *The Object/Field Divide*: it has been recognized that objects and fields – the assignment of values to spatial locations – have to coexist in geographic applications [11]. However, this distinction still constitutes a problem for the object representation tradition. Indeed, *why* is an object such as a mountain a field or an object, or, better, *when* do we want it to be a field or an object? What about fields (e.g. demographics) which are composed of individual objects? What about fields composed of other fields (e.g. land coverage)? Human cognition never fails in choosing the best representation, object or field or composition of both, according to a task and a context.
- *The Cognitive Imperative*: space is *experienced before being known*, as shown by Naïve Geography [3], which shows to what extent useful representations of space are to be rooted into human cognition. This is demonstrated in a different way by new web maps, in which multiple reality effects are embedded, such as seamless continuity in map browsing instead of image by image retrieval, satellite imagery, road level or oblique photography, 2.5 or even 3D features. These representations are appealing since they allow leaving at will the world of iconic or symbolic representation to apply cognitive models we use commonly in our daily life. Examples of these models are *affordances* [13] (what an element of the external world *allows me to do* is more important than its other characteristics), *image schemata* [5] (an element can be further reduced to simple concepts which are self-understandable), or *conceptual spaces* [14] (a concept is a point in a multi dimensional space of simpler representations).
- *The Multi-Representation Problem*: beyond the fundamental object field divide, or the cognitive approach to object representation, a geographical object simply

changes according to the level of detail needed or requested, and to the task at hand. For example an airport will be a node in a flights graph from an international point of view, then become an independent region in a land cover study, or a simple traffic node, or a complex environment itself containing a road network and buildings, or a group of 3D structures with emergency access path in a fire escape scenario, etc. The multitude of contexts and corresponding relevant representations raises the question of the possible uniqueness of geographic object representation; indeed, if many representations are useful how can they be linked and accessed in a timely manner, according to contextual information?

The e-Merges application aims to eventually address these concerns, by linking them via the notion of *spatial context*. Indeed, in order to ultimately (a) alternate object and field representations, (b) provide cognitively relevant information, and (c) choose between multiple representations of the same element, the representation of *spatial objects* becomes *spatial context* dependent. We are going to define both notions in turn.

Spatial Objects

In order to describe and to reason about Geographic Objects in all their generality [3] a simple yet precise definition is needed. Our model is based on Galton's theory of objects and fields [6], although, following a pragmatic approach, we only used the aspects useful to the e-Merges prototype.

In this approach an entity inheriting from the *spatial-object* concept simply acquires a location, which is now restricted to a polygonal area. This concept is used to provide a standard spatial representation to other entities. Mapping of arbitrary domain objects to the spatial ontology can be *automatic* or *manual*. In automatic mapping, a procedure collects each object's attribute value and transforms it into an attribute name/value pair of a spatial object. In manual mapping, any transformation becomes possible.

To achieve separation of concerns, two other ontologies are used. The *Archetypes* ontology provides very high level abstractions (e.g. *container*, *house*, *agent,* etc.) to which entities have to be mapped. In this way even if the client application does not understand the type of element that is to be represented, displaying them is still possible by using the attached archetype, which clients are requested to be familiar with. For example if a client application does not have a representation for a *hospital*, it would now how to represent a *house* as the attached archetype, which is in any case more sound than other archetypal representations such as *agent* or *link*. To this ontology can be attached *image schematic* features allowing standard comportments to be displayed.

Finally, the *HCI* ontology is a view of an object as it is to be displayed in a user interface. For example some interfaces allow information to be displayed when hovering with the mouse over an object; an attribute of an adapted HCI concept allows us to specify which information, by automatic mapping (e.g. a procedure choosing any slot containing the string "id" or "name") or with a manual one.

These ontologies, together with the attached mapping mechanisms, are called *integration ontologies* since they allow the integration of spatially related data sources ranging over very different domains. However, as the number of data sources increases, the task of presenting objects and possible queries according to the context, in order for the user not to be overwhelmed by the amount of data and services, becomes essential. The notion of spatial context is used to provide only relevant information and services.

Spatial Context

In order to alternate representations, and to stay cognitively sound, an application needs to be *context-aware* [12]. In the context of GIS we believe that context information is mainly related to the *user*, the *task*, the *location*, and the *focus* of interest. Indeed a user first identifies him- or herself, to achieve a task that is defined by the first actions achieved, which are relevant in a precise area; furthermore if the action has consequences, e.g. if new spatial objects are retrieved, focus on an element is an indication of the user's intent, and can be revealed by a click on this element.

Several elements may change according to the context. Firstly, when there is room for change, object representations differ according to the context; e.g. the town council may get a representation of an area showing parcel ownership, while the fire brigade may get access roads and water points. Secondly, to objects and to situations are linked goals, which allow getting more information in a precise context. For example an area defined as an evacuation zone may offer goals allowing finding the nearest supermarkets or hotels, etc. This links the SWS notion of *goal* to the cognitive notion of *affordances* attached to an object.

The question of whether a specific context reasoning engine has to be used is open. However, we believe that in the context of SWS, a more scalable solution may be achieved by distributing the task of context handling amongst *smart services* which also implement reasoning in our architecture. Indeed context pervades the elements of an application, and can be represented (a) at a goal level, i.e. by offering very specific goals only, according to the context, e.g. a *get-heated-shelters* goal will be presented in an emergency case involving low temperatures, or (b) at a composition level, i.e. generic goals are presented and smart composition ensures context relevance, e.g. the generic goal *get-shelters* is presented to the user but highlights heated shelters according to the task. The first solution has the advantage of being more explicit, whilst the second is easier to implement since it requires fewer goal definitions. Being able to handle context at every level makes both solutions possible in SWS based applications.

The IRS-III Approach to SWS Applications

Applications using IRS-III follow a layered approach (cf. Fig. 1) in which (micro-) functionalities of *legacy systems* are exposed through *Web Services* – based on standards or on REST – and described with ontologies. These *Semantic Web Services* can then be invoked from the (web) *presentation* layer, by using a provided API, SOAP messages, or the REST protocol.

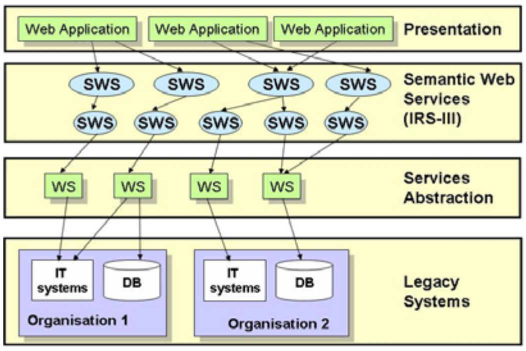

Fig. 1 The generic architecture used when creating IRS-III based applications.

The *Web Service Modelling Ontology (WSMO)* is a formal ontology for describing the various aspects of services to enable the automation of WS discovery, composition, mediation and invocation. The meta-model of WSMO defines four top level elements: *Ontologies*, *Goals*, *Web Services*, and *Mediators*.

Ontologies [2] provide the foundation for describing domains semantically. They are used by the three other WSMO components. *Goals* define the tasks that a service requester expects WSs to fulfil. In this sense they tend to reflect the service user's intent. *Web Service* descriptions represent, in terms of *capabilities* (what the service can do) and *interface* (how to use it), the behaviour of a deployed Web Service. The description also indicates how WS communicate (*choreography*) and how they are composed (*orchestration*). *Mediators* handle issues of data and process interoperability that arise between heterogeneous systems. One of the characterizing features of WSMO is that all components – Ontologies, Goals and Web Services – are linked by Mediators. In particular, WSMO provides four kinds of mediators:

- *oo-mediators* for mediating between heterogeneous ontologies;
- *ww-mediators* connect WS to WS;
- *wg-mediators* connect WS with Goals;
- *gg-mediators* link different Goals, solving input conflicts and transforming processes.

By extending WSMO's goal and Web Service concepts, clients of IRS-III can invoke web services via goals. That is, IRS-III supports so called *capability-*, or *goal-driven* service invocation which allows the user to use only generic inputs, hiding the possible complexity of a chain of heterogeneous WS invocations. The decoupling of the actual user vision of a task and its execution allows us to get closer to the user's cognition of the situation and task. Mediators link goal and web services, solving existing mismatches, and allowing complex composition of services to be constructed.

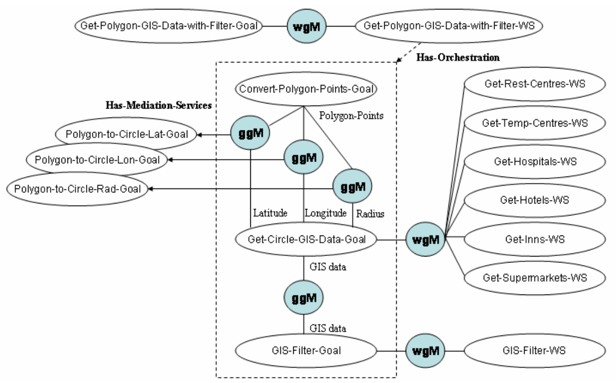

Fig. 2 Structure of the WSMO description of the e-Merges prototype. To avoid cluttering the diagram, wgM and Web Services balloons were omitted.

To illustrate such a composition we describe in the following (Fig. 2) the structure of the WSMO descriptions associated with an example goal, *Get-Polygon-GIS-data-with-Filter-Goal*. This goal describes the request of a class of shelter (hospital, inn, hotel, etc.) in a delimited query area. The user selects a class of shelters while the *polygon* query area is interpreted through context. However, the WS at hand only returns a specific class of shelter in a *circular* query area. The results also have to be filtered in order to return only shelter relevant to the task (in our case, the management of a snowstorm emergency). The problems are: (1) *selection* of the adequate WS; (2) *mediation* of the different area representations (polygon *vs.* circular); (3) *orchestration* of the retrieve and filter data operations. IRS-III offers approaches to solve these problems:

- *WS Selection*: each WSMO description of WS defines, in its *capability*, the specific class of shelter that the service provides. All descriptions are linked to *Get-Circle-GIS-Data-Goal* by means of a unique *wg-mediator (wgM)*. The goal expects as input a class of shelter, and a circular query area. At invocation time IRS-III discovers through the wgM the WS associated to it. Then it selects one amongst them according to the specific class of shelter described in WS capabilities.

- *Area mediation and orchestration*: *Get-Polygon-GIS-data-with-Filter-Goal* is associated to a unique web service that orchestrates – here, invokes in sequence – three sub-goals. The first one simply gets the list of polygon edges from the input; the second is the above mentioned *Get-Circle-GIS-Data-Goal*; and finally the third invokes the smart service that filters the list of GIS data. The first two sub-goals are linked by means of three *gg-mediators (ggM)* that convert the list of polygon edges provided by the first sub-goal to the centre (latitude and longitude) and radius of the circle that circumscribes that polygon. To accomplish this, we created three mediation services invoked through *Polygon-to-Circle-Lat-Goal*, *Polygon-to-Circle-Lon-Goal*, and *Polygon-to-Circle-Rad-Goal*. The results of the mediation services and the class of shelter are the inputs of the second sub-goal. A unique ggM connects the output of the second to the input of the third sub-goal. No mediation service is necessary here.

The e-Merges Prototype

The prototype was designed for the *Essex County Council (ECC) Emergency Planning Department*. The ECC is a large local authority in South-East England (UK). Following several interviews with spatial data holders in the ECC it was decided to focus the scenario on the ECC Emergency Planning department, and precisely, on a previous emergency situation: the snowstorm which occurred in the vicinity of Stansted airport on the 31st of January 2003. To avoid interferences, data from the ECC Emergency Department and the Meteorological Office was replicated. This will also allow us to compare emergency officer's decisions regarding contact with rescue corps and voluntary associations, or actions necessary to provide refuge and supplies to trapped travelers, etc. – with those of the prototype users.

The e-Merges prototype is a decision support system, which assists the *Emergency Officer (EO)* in handling the dynamics of an emergency situation and gathering information related to a certain type of event, faster and with increased precision.

Data was integrated from three different sources. UK's *Meteorological Office* providing snow level information, the *ViewEssex* database, a centralized database maintained by British Telecommunications (BT) managing spatial-data for the ECC, and *BuddySpace*, an Instant Messaging client built on top of the *Jabber*[4] protocol and providing lightweight communication and collaboration means [4].

A Web interface based on *Google Maps* supports the spatial representation part of the applications. This interface is web standards based, using *XHTML* and *css* for presentation. *JavaScript* is used to handle user interaction as well as *AJAX* techniques for IRS-III goal invocation. The significant components of the interface are a central *map*, which uses the Google Maps API to display *polygons* and *objects* (custom images) at specific coordinates and zoom levels. Context is manifest in that objects have corresponding *goals* and *attributes*, which are displayed in a pop up window or in a hovering transparent region above it.

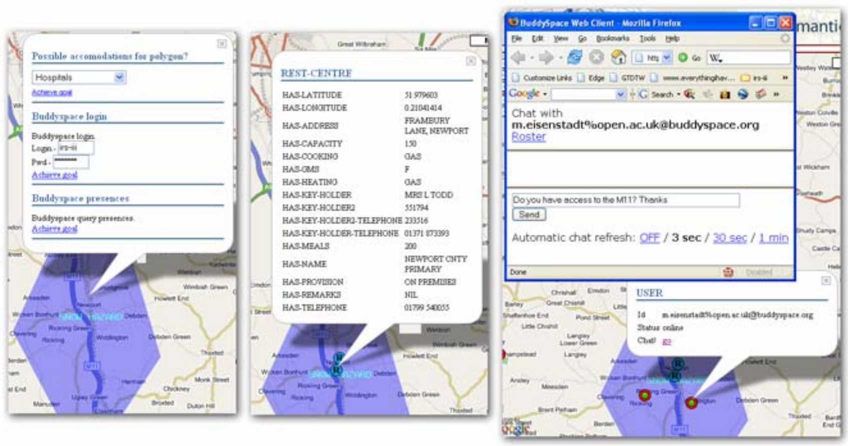

Fig. 3 Once defined the area presents goals which can be queried to obtain objects and allow further interaction.

---

[4] http://www.jabber.org/

As an example of practical usage, we describe how an EO gathers information regarding an emergency situation (a *snow hazard* or a *snow storm* each offering different goals), before trying to contact relevant agents. The procedure is as follows:

1. Based on weather information the EO draws a polygon on the map and assigns a type of emergency to the region. Here, *a snow storm*.
2. Described in an ontology, the new instance has attached *attributes* and *goals*. Three goals are attached to the emergency, one *gets shelters* at distance from the area, two others *connect* to BuddySpace and *get relevant presences*. (Fig. 3 left)
3. First, the user requests all *rest centres* inside the region. They are retrieved with their features and attached goals. (Fig. 3 middle)
4. With that information the EO *logs* into BuddySpace, then *contacts* the relevant persons to request action or information. (Fig. 3 right)

A screencast of the interaction as well as a live version are available online[5], to be used with the latest version of the Firefox Web browser[6].

The integration of new data sources is relatively simple although not totally trivial. Indeed IRS-III SWS integration allows the description of *any* data source available on the web, whilst the application aim is to represent it in a definite and dynamic *context* (described in an ontology), unlike mashup builders[7], which only gather syntactically similar feeds on a map. The steps involved in the process of adding a new data source, as well as the ability to automate them, are described in the following:

- *Ontological description of service*: the service, composed of the data types involved as well as its interface, has to be described in a low level ontology, usually at a low enough level to remain close to the data. This step can be automated in many cases.
- *Lifting definition*: the lifting operation allows the passage of data type instances from a syntactic level (XML) defined in the data schema to an ontological one (OCML) specified in the ontology definition. This process can be automated every time the previous step can be.
- *Mapping to integration ontologies*: this process can be fully automated by default, and customized as needed.
- *Goal description*: a new goal has to be defined which represents the newly integrated web service.
- *Mediator description*: the goal has to be linked to the WS with a mediator, which is often a trivial operation.
- *Lowering definition*: the lowering operation transforms instances of aggregation ontologies into syntactic documents to be used by the server and client applications.

---

[5] http://irs-test.open.ac.uk/sgis-dev/
[6] http://www.mozilla.com/firefox/
[7] e.g. http://mapufacture.com/georss/

Conclusion

The e-Merges approach to spatial data integration presents advantages from an end user as well as from a knowledge expert point of view. Indeed it allows the end user to handle tasks in a data rich environment without being overwhelmed by the amount of information and the complexity of queries. From a knowledge expert point of view the data source integration approach presents many advantages compared to standard based approaches such as the one demonstrated in the OWS-3 Initiative[8], including framework openness (i.e. standards make integration easier but are not mandatory) and high level service support (i.e. all the benefits of the underlying SWS platform, such as discovery, composition, etc. are immediately available).

Future developments will include an increase in the complexity of the integration ontologies (spatial, HCI and archetypes) in order to allow fields, multi representation and cognitive features to be manifested in the interface. Also, making the integration of new data sources even easier constitutes a long term goal for the IRS-III SWS execution platform.

References

1. Domingue, L. et al., IRS-III: A Platform and Infrastructure for Creating WSMO-based Semantic Web Services. (2004)
2. Gruber, T. R., A Translation Approach to Portable Ontology Specifications. (1993)
3. Egenhofer, M. J. & Mark, D. M., Naive geography. (1995)
4. Eisenstadt, M., et al., Instant messaging + maps = powerful collaboration tools for distance learning. (2003)
5. Mark, D., Cognitive Image-Schemata for Geographic Information. (1989)
6. Galton. A., A formal theory of objects and fields. (2001)
7. Motta, E., An Overview of the OCML Modelling Language. (1998)
8. Roman, D., Web service modeling ontology standard (WSMO-standard). (2004)
9. Peuquet, D., Representations of Space and Time. (2002)
10. Smith, B. & Mark, D., Ontology with Human Subjects Testing. (1999)
11. Couclelis, H., People Manipulate Objects (but Cultivate Fields). (1992)
12. Dey, A. K. & Abowd, G. D., Towards a Better Understanding of Context and Context-Awareness. (1999)
13. Gibson, J. J., The theory of affordances. (1977)
14. Gärdenfors, P., Conceptual Spaces. (2000)

[8] http://www.opengeospatial.org/initiatives/?iid=162