

Ontology Driven Visualisation of Maps with SVG

An Example for Semantic Programming

Frank Ipfelkofer, Bernhard Lorenz and Hans Jürgen Ohlbach
Institute for Informatics, Ludwig-Maximilians University, Munich,
www.pms.ifi.lmu.de, ipfelkofer@kastner.de, {lorenz,ohlbach}@pms.ifi.lmu.de

Abstract

This article describes a particular use of ontologies for visualising maps in a browser window. Geographic data are represented as instances of concepts in an ontology of transportation networks which was designed in close relation to GDF. These data are transformed into SVG, whereas the transformation is specified symbolically as instances of a transformation ontology. This approach is extremely flexible and easily extendible to include all kinds of information in the generated maps. The basic implementation technique is to use classes and instances of ontologies in an intelligent way. This article is an abbreviated version of the research report [12].

Keywords— semantic web, geospatial notions, ontologies, visualisation, scalable vector graphics, semantic techniques

1 Introduction

There are many different ways for generating maps as images on a computer. The most straightforward way is to read the geographic data from a data source and to use special purpose algorithms that transform the data into some bitmap graphics format. These algorithms are rather complex and not easy to change or extend. Only experts who are familiar with the details of the process can do this. The algorithms depend very much on the particular data format and they usually yield only static pictures.

An alternative method is to generate output by means of ontologies and ontology instances instead of using specialised algorithms processing data transformed into generalised formats. Furthermore, instead of creating bitmap graphics, the results are encoded in a graphics description language, such as SVG, which is used as an example throughout this article. Several advantages result from this approach which are briefly laid out in this introduction and discussed in more detail in section 2.

Scalable Vector Graphics (SVG) [19, 1] is an XML-based language for describing geometric objects. There are special plugins for web browsers which can render SVG files in a browser window [1]. Compared to bitmap graphics, SVG has a number of advantages:

- SVG is based on vector graphics, which are zoomable without losing resolution on the screen.
- SVG has language constructs for describing dynamic changes of the graphics.
- Since SVG objects have a DOM representation [4] in the browser, script languages like JavaScript can interact with it. SVG documents can therefore serve as GUIs to interact with the user. We used this to allow the user to interactively change the presentation of the maps.
- SVG renderers can adapt the generated picture to the output device. Therefore the SVG generator need not worry about the device characteristics.
- There are different SVG sublanguages available [18, 20] to provide adaptation means for very different output media.
- SVG documents are XML documents which can be read by humans. This is very useful during the development and test phase of SVG generators.

The generation of code in a graphics description language like SVG instead of bitmap graphics therefore has the following advantages:

- The algorithms for transforming the geographic data are much simpler because the final rendering of the graphical data is done by the browser.
- The generated graphics code is device independent. The renderer automatically adapts the graphics to the output device.
- If the graphics description language has constructs for dynamic presentation, they can be used directly without having to care about rendering issues.
- Most kinds of interaction, for example zooming in and out to a certain extent, is done by the browser, and need not be taken into account by the transformation algorithms.

The primary data sources for the visualisation are usually Geographical Information System (GIS) databases which (in-)directly provide data in standard formats, for example the Geographic Data Format (GDF) [6] or the Geography Markup Language (GML) [7]. This is not the only choice. In this article we propose an alternative. We still use GIS data in some of the standard formats as primary source, but only because these are the only available data. The idea is to take an OWL ontology of transportation networks to represent the data as instances of the concepts of this ontology. OWL provides a data format for instances of the concepts of the ontology (basically the Resource Description Framework (RDF)), and we use RDF for the GIS data. This is not just a syntactic reformulation. It offers completely new possibilities because the data format is only loosely coupled with the ontology. For example, consider an ontology containing the concept of a *road*. A road may have directions, at most two. If there is a particular road *R* with directions = 1 then OWL would classify *R* as a *road*. If, in a later step, the ontology is extended with the concept *one_way_road* as a *road* with directions = 1 then OWL would automatically reclassify *R* as a *one_way_road*. Thus, there is a certain degree of independence between the OWL data format and the ontology. The same data can be used for different ontologies.

We developed the Ontology of Transportation Networks (OTN) [2]. OTN was generated by making the concepts and structures which are implicitly contained in GDF [5, 6] explicit as an OWL ontology. OTN contains all kinds of notions for transportation networks, roads, trains, ferries and much more.

The method for visualising maps uses three classes of input data:

1. Concrete GIS data which has been transformed into the OWL data format (we used the map of the city of Munich).
2. An ontology of transportation networks, in our case OTN.
3. A set of transformation rules which determine how the instances of the concepts in the ontology are to be transformed into SVG code.

The transformation algorithm now applies the transformation rules to all relevant instances of the concepts in the ontology and produces SVG documents as output. With this architecture it is extremely easy to change the visualisation. For example, if we want to distinguish one-way roads from ordinary roads, it is only necessary to introduce the concept of a one-way road in the ontology and to add a corresponding transformation rule.

If the data source is in XML, which is the case for GML or the OWL data format, and the target is also XML, which

is the case for SVG, there is a further alternative for generating maps. One can develop an XSLT style sheet that transforms the GIS data into SVG [14]. This approach is, however, very limited and relatively inflexible because the XSLT style sheet depends extremely on the structure of the XML data source. Moreover, it does not support rasterisation and levels of detail, which is extremely important for working with large maps.

In this document we describe the basic ideas and techniques of our approach, a more detailed description can be found in [11]. Section 2 describes the limits of SVG and illustrates our solutions, particularly regarding a dynamic loading mechanism and a special rasterisation technique based on R-trees. Detailed information about the transformation step can be found in [12]. While section 3 briefly indicates a number of possibilities to extend the system and its functionalities, related work is shown in section 4. We conclude this paper with the summary and outlook in section 5.

2 SVG Visualisation

We start with a description of the SVG visualisation technique because this motivates some of the design decisions for the transformation method.

2.1 The Final Result

The final result of the visualisation is illustrated in fig. 1. The browser window consists of a frame containing the SVG map and a HTML menu on the right hand side. The SVG map is zoomable in a wide range, more than the built-in SVG zooming facility would allow. The section to be displayed can be changed by just dragging the mouse over the window. The map may contain dynamic elements (e.g. buses or trains moving along the rails, clouds moving over the scene, etc.).

The menu allows the user to choose what he wants to see. It is divided into two main sections, *Modules* and *Ontology*. The *Ontology* section corresponds to a *display ontology*, which is a tree of concepts in the transportation network realm. Checking or unchecking the appropriate box causes the corresponding items in the map to become visible or invisible. A *module* in the upper section of the menu consists of a set of elements from the display ontology. Any combination of elements from the display ontology can form a module. Checking or unchecking makes the whole group of items visible or invisible.

There is a further feature which is not part of the section of the browser window shown in fig. 1. Below the map there is a text input section where one can type in a street name and the street is then highlighted in the map.

2.2 Dynamic Loading

All visualisation systems for maps have the same problem: the server has usually much more data than the user at

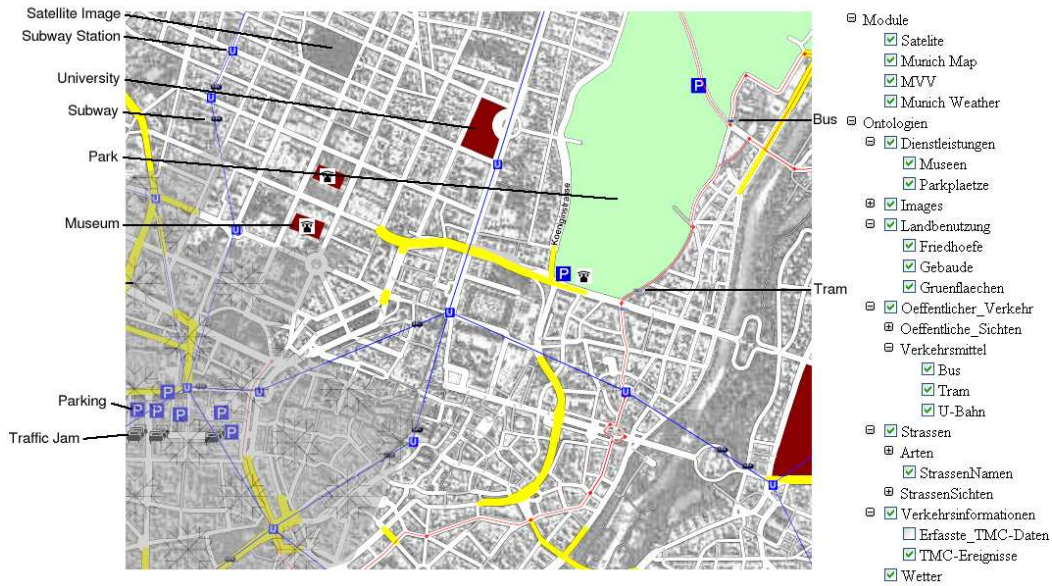


Figure 1: Visualisation in SVG

the client side wants to see. Network bandwidth and computation capacity are not large enough to transfer all the data from the server to the client such that the client can decide what to show to the user and what not. Therefore it is necessary to partition the data at the server side and send only the relevant parts to the client side. If the user changes the section of the map to be shown or he zooms in and out, more data needs to be loaded dynamically during the user interaction.

SVG has no built-in facility for dynamically loading data from the server. The combination of the DOM representation of SVG data in the browser memory and the possibilities of scripting languages like JavaScript to modify the DOM at any time, however, makes it possible to program dynamic downloading of data from the server (for a detailed description see [12]).

In order to use this mechanism for loading only the actually needed parts of a big map, we have to solve two problems. The first problem is to divide a big map into small enough tiles which can be loaded independently. The second problem is to support zooming already at the server side. The problem here is that the same item must be displayed differently at different zoom levels. For example, if the whole map of Germany is to be shown, it makes no sense to display all the details of, say, the city of Munich. Munich should in this case be displayed only as a dot, or maybe as a very simple polygon filled with a uniform colour. If the user zooms into Munich, he wants to

see of course more detail. These different views have to be prepared at the server side.

2.3 Rasterisation

There is a very simple solution for splitting a big map into smaller tiles: a fixed grid is imposed on the map and the map is split into the grid elements. The disadvantage is that this way the split parts may have very different size – we are talking about vector data here. There may be split parts with almost nothing in it and split parts in very densely populated areas which contain thousands or millions of items. One could argue that this is not the case with rasterised data, since all tiles would come in the same size (e.g. 128x128 pixels). The contained information, the net data, however, would also greatly differ.

A much better distribution of equally sized split parts can be obtained with R-Trees [16, 17]. An R-Tree is a structure for storing 2-dimensional data. Each node in the tree contains data about its *minimum bounding rectangle*, i.e. the smallest rectangle which includes the node itself and all of its descendants. Leaves contain the data and nodes contain index information. Nodes can be further combined within other nodes, rectangles can be overlapping. The root node therefore contains all descending nodes and subsequently all leaves including the bounding rectangle of the whole tree. It is normally not possible to generate an optimal R-Tree, since this would involve a complexity of $O(n) = 2^n$. Therefore, there are generally a number of different (equal) instances of an R-Tree and some algorithms

for maintaining its structure. This does, however, not have a significant impact on the rasterising process as we need it for our application. Fig. 2 shows a possible R-Tree, which is rather self explanatory.




Figure 2: R-Tree Sample

In our special case, using an R-Tree means using the advantages of rectangular grid sections while eliminating the need for separate indexing or cumbersome preprocessing of data. The grid is comprised of the minimum bounding rectangles (which can be overlapping), whereas each element belongs to only one section and all sections contain a similar number of elements. Each node contains information about its children and the sizes of their respective rectangles. This allows for recursive search from the root along the different nodes, while for each node it can be quickly decided, whether it touches the area to be displayed (and therefore, whether data from its children has to be loaded). Elements can easily be distributed equally between grid sections and there is no need for a separate index file. Fig. 3 shows a tile of a map of Munich which has been generated using an R-Tree. This tile contains only data for a particular road type. We use the OTN ontology to separate the items in the tiles into instances of the same classes.



Figure 3: A generated tile of a map

2.4 Levels of Detail

Yet another similar problem stems from the zoom mechanisms. Depending on the current zoom level, certain elements - mostly because of their size - cannot be displayed

properly because they would be too small to be useful. Elements like this include smaller streets (which come in greater numbers as well), street names and similar things. Moreover it makes sense to simplify certain elements to simpler structures in order to improve readability and usability of the map. Cities might be reduced to circles or dots of different diameter (depending on other attributes, such as number of inhabitants). This form of presentation is much more useful than putting processing power into rendering irregular city boundaries which are too small to be identified as such.

Every element therefore contains the attributes *minDetail* and *maxDetail* which set the levels of detail between which the element is to be visible. The level of detail is the minimum of both the vertical and horizontal resolution. On a map displaying from coordinates (200, 400) to (500, 1000) the minimum would be $\min(300, 600) = 300$.

Our system therefore generates SVG files for the different tiles at the different zoom levels. If the user zooms in or out the corresponding files are automatically loaded.

3 Further Services

The browser that renders SVG data has, via JavaScript, access to the data structures underlying the items on the generated image. This can be exploited to implement further services. One of the services we implemented is a road finder. Since every road has a name, the browser can build an index for the roads when they are downloaded. The user can now type in the name of a road and the browser uses the index to match the road name with the SVG elements that display the road. These elements are now highlighted by changing their colour attribute. So far this works only for roads. The reason is that different types of objects are usually represented by different combinations of SVG elements. A road, for example, consists of road segments which are displayed with one or two (or more) SVG path elements. This association is different for other types of objects and has to be programmed in another way.

Highlighting particular roads is a service which can be executed at client side. We also implemented a prototype of a service where the client has to contact the server. This service searches the shortest path between two locations on a map. The client sends the two locations to the server, the server computes the shortest path, and sends back an SVG document that shows the shortest path in the browser window. So far, only the interface between client and server is implemented and only fixed test data are sent.

4 Alternative Approaches

As mentioned before, there exist a number of different possibilities in order to provide similar services. We give some examples here which use different approaches, although none incorporate a holistic use of ontologies.

Two of the most prominent commercial examples come from the search engine provider Google. *Google Maps* and *Google Earth* show two very different approaches in client-side applications for GIS data presentation, a more comprehensive description can be found in [21, 22]. Google Maps is based on JavaScript and XML, and can be accessed with any current web browser, whereas Google Earth is a proprietary stand-alone application.

Google Maps is a free Web Map Server application [9] which provides zoomable and pannable street map and satellite images for the whole planet, along with route planning and business locator facilities for a number of countries¹. Following Google's key mission, Google Maps can be combined with some search functionality. Search results can for example be restricted to a certain area: "Pizza in Boston" yields facilities providing pizza in the greater Boston area. This applies to other services as well. Further functionality includes common routing and navigation, including lists of driving directions².

Although in the early stages the underlying protocols and mechanisms have not been publicly available, reverse engineering of the interface (which is mainly based on JavaScript and XML) has led to the development of expanded and customised features. Using the core engine and the map/satellite images hosted by Google, such expansions can introduce custom location icons, location coordinates and meta data, and even custom map image sources (e.g. [10, 15]). Meanwhile Google released a Google Maps API to Google developers for non-commercial purposes.

Google Earth, formerly developed as a purely commercial product by Keyhole Inc. and now owned and made freely available by Google, is a virtual globe enabling the viewing of vectorised and rasterised data by generating views of the earth from above [8]. Currently Google Earth is only running on personal computers using Microsoft Windows, although versions for Linux and Mac OS versions have been announced for the end of 2005.

Google Earth operates in a similar manner to Google Maps, but as a 3-dimensional application. Instead of planar maps, Google Earth provides a globe, which can be viewed, rotated, zoomed into, much like its real life counterpart. The most important difference is a layer architecture, well known from Geographic Information Systems, which contain different sets of features, such as parks, rivers, roads, borders of countries, locations of national monuments and many thousands of other places. These are provided not only by Google, but by the whole Internet community. These layers can be selected and deselected by the user in order to create a view containing the desired features. The mechanisms for incorporating cus-

tomised data into the client are publicly available via the Keyhole Markup Language (KML) [13].

MacauMap[3] is a handheld digital map application which displays information about tourist-related spots (hotels, restaurants, etc.), provides a bus routing function for calculating an optimal bus route between a pair of bus stops and offers other functions. As it is targeted for mobile use, the focus of development mainly lies on two issues pertaining to mobile devices: device resources and user interface. Since computing power, bandwidth and memory capacity are restricted, special techniques have been developed to optimise data processing. Likewise, user interaction is restricted to either a stylus (PDAs) or an alphanumeric keypad (mobile phones) the user interface has been adapted accordingly.

5 Summary and Outlook

In this work we illustrate a particular use of ontologies for dealing with geographic data. The geographic data are represented in the OWL data format that corresponds to the Ontology of Transportation Networks (OTN). Since there is a certain degree of independence between the data and the ontology, it is possible to adapt the ontology to the needs of the application and still work with the same data.

The transformation of the geographic data into SVG is also controlled by an ontology. The SVG elements are represented as concepts of a *transformation ontology* and the particular rules for transforming the data in a particular way are specified as instances of the concepts of the transformation ontology. By changing these instances or creating new instances one can change or extend the displayed maps very easily. Therefore this is an extremely flexible architecture which allows one to program the generation of maps by specifying the transformation in a symbolic way.

There are two main differences to other approaches for generating maps:

1. the programming technique is *semantic* with a significant symbolic specification part. This is much more flexible than other programming techniques. For example, if the presentation of the map is to be changed or extended, it is usually sufficient to start an OWL editor and change some classes or instances;
2. behind the displayed map there is always the document object model (DOM), and the elements of the DOM are still linked to the ontology. This enables interactive services where the ontology can be invoked.

¹The U.S., Canada, Japan, Hong Kong, China, the UK and Ireland (city centres only)

²As of June 2005, Google Maps features road maps for the United States, Puerto Rico, Canada, and the United Kingdom.

SVG is a very expressive language with a number of quite powerful features. The renderers for SVG are therefore quite complicated and still seem not to be optimised for large data sets. The dynamic features of SVG in particular slow down the rendering considerably. Rendering big maps with a lot of items is still slower than it could be. We are currently exploring the possibility to write a renderer in Java, not for full SVG, but for the special constructs needed for visualising maps. This should give similar performance as for example map24 or the aforementioned Google Maps.

In contrast to the commercial examples described in section 4 we didn't develop yet another set of proprietary languages, interfaces and methods, but instead use already available open standards. Apart from the obvious advantages, this makes adaptation to different platforms and devices much easier.

The techniques can also be extended in other directions. For example, one could integrate scrollbars which are to be used to make the graphics dependent on further parameters, in particular time. This way, time dependent data can be integrated into the graphics in such a way that the user can choose the time and the graphics is automatically adapted. Another direction could be to replace SVG with X3D in order to present 3D graphics. The transformation techniques are similar to the 2D case.

Acknowledgement

This research has been co-funded by the European Commission and by the Swiss Federal Office for Education and Science within the 6th Framework Programme project REWERSE number 506779 (cf. <http://rewerse.net>).

References

- [1] Adobe SVG Viewer, Version 3.01. <http://www.adobe.com/svg/>, September 2003.
- [2] Bernhard Lorenz and Hans Jürgen Ohlbach and Laibing Yang. *Ontology of Transportation Networks*. REWERSE Deliverable A1-D4, University of Munich, Institute for Informatics, 2005.
- [3] Robert P. Biuk-Aghai. *MacauMap: Next Generation Mobile Travelling Assistant*. In *Proceedings Of Map Asia 2004, Beijing, China*, August 2004.
- [4] Document Object Model (DOM) Level 2 Core Specification. <http://www.w3.org/TR/DOM-Level-2-Core>, November 2000.
- [5] International Organisation for Standardisation (ISO). *Geographic Data Files 3.0 (GDF) Documentation*. http://www.ertico.com/en/links/links/gdf_-_geographic_data_files.htm, 1995.
- [6] International Organisation for Standardisation (ISO). *Intelligent transport systems - Geographic Data Files 4.0 (GDF) - Overall data specification*, ISO/DIS 14825/2004, February 2004.
- [7] Geography Markup Language GML, Version 3. <http://www.opengis.org/docs/02-023r4.pdf>, (accessed 11/2005).
- [8] Google Earth. <http://earth.google.com>, (accessed 11/2005).
- [9] Google Maps. <http://maps.google.com>, (accessed 11/2005).
- [10] Greasemonkey. <http://greasemonkey.mozdev.org>, (accessed 11/2005).
- [11] Frank Ipfelkofer. *Basisontologie und Anwendungs-Framework für Visualisierung und Geospatial Reasoning*. Diploma thesis, University of Munich, Institute for Informatics, 2004.
- [12] Frank Ipfelkofer, Bernhard Lorenz, and Hans Jürgen Ohlbach. *Ontology Driven Visualisation of Maps with SVG – An Example for Semantic Programming*. Forschungsbericht/research report PMS-FB-2006-5, Institute for Informatics, University of Munich, 2006.
- [13] Google Earth KML Tutorial. http://www.keyhole.com/kml/kml_tut.html.
- [14] Andreas Kupfer. *Visualisierung von GML mit XSLT und SVG*. Diploma thesis, Technical University of Braunschweig, 2003.
- [15] MyGMaps. <http://mygmaps.com/mygmaps.cgi>, (accessed 11/2005).
- [16] R-Tree Portal. <http://www.rtreeportal.org>, Juni 2003.
- [17] R-Tree Visualization Demo, National Technical University of Athens. <http://www.dbnet.ece.ntua.gr/~mario/rtree>, November 1999.
- [18] Mobile SVG Profiles: SVG Tiny and SVG Basic, W3C Recommendation. <http://www.w3.org/TR/SVGMobile>, January 2003.
- [19] Scalable Vector Graphics (SVG) 1.1 Specification, W3C Recommendation. <http://www.w3.org/TR/SVG11>, January 2003.
- [20] SVG Print, W3C Working Draft. <http://www.w3.org/TR/SVGPrint/>, July 2003.
- [21] Wikipedia: Google Earth. http://en.wikipedia.org/wiki/Google_Earth.
- [22] Wikipedia: Google Maps. http://en.wikipedia.org/wiki/Google_Maps.