Local Spatial Data Infrastructures Based on a Service-Oriented Architecture

Clodoveu A. Davis Jr.[1], Leonardo Lacerda Alves[1]

[1]Instituto de Informática – Pontifícia Universidade Católica de Minas Gerais
Rua Walter Ianni, 255 – 31980-110 – Belo Horizonte – MG – Brazil

clodoveu@pucminas.br, leonardo@lacerda.eti.br

Abstract. *Sharing geographic information is an essential activity which has been sought since the early days of GIS, mostly due to the cost of information collection and maintenance. Having once depended on the establishment of data transfer standards, sharing initiatives gradually evolved towards the creation of clearinghouses, Web resources that centralize links to various GI sources, but are still data-oriented. The current focus on spatial data infrastructures changes that, by establishing a service-oriented view, thus allowing for the creation of shared, distributed, and interoperable environments through Web services. This paper explores, in a preliminary fashion, such an architecture as applied to distributed geographic applications, focusing on the potential for local services and local uses, and proposing specialized services deemed essential for urban-scale applications.*

1. Introduction

The problem of sharing large volumes of spatial information is becoming increasingly important. Geographic information systems (GIS) have evolved from a project-oriented tool to become managers of enterprise-wide information resources. Societal aspects of geographic information have also emerged as important resources for governmental actions. With this, GIS are increasingly becoming the core of computational environments involving large numbers of users, spread over numerous locations, and high volumes of data (INSPIRE Architecture and Standards Working Group 2002).

One of the greatest of the early promises of GIS, namely the establishment of geographic location as a means to link data sets from apparently disparate origins , makes much more sense if we consider a computational environment in which data are freely shared in a seamless way. This is not the current practice. Organizations that intend to share information among themselves usually face considerable challenges, regarding data storage formats, data quality issues, content issues, cartographic projection parameters, and even data structures (e.g., object-oriented versus topology-oriented) (Rajabifard and Williamson 2001). A multitude of data translation tools has become available to partially solve some of those problems, both as part of GIS packages or as independent, user-developed applications (see, for instance, the "data translation" section of the GeoCommunity Web site[1]).

[1] www.geocomm.com

Thus, GIS communities have been experiencing a gradual shift from a past in which no integration was possible, to a present in which offline data replication and translation is the norm. As to the future, we notice that semantic interoperability is still far from existing tools, even though its promise generated the benefit of more interest in metadata (FGDC 2001) and in ontologies (Fonseca, Egenhofer et al. 2000) throughout GIS communities. Nevertheless, data exchange arrangements increasingly claim for systems development alternatives in which data from external providers can be used without translation, and without the updating delays that result from offline usage.

This paper presents a novel architecture for the development of geographic applications, in which data are provided as a number of distinct network-based information services, thus forming a *spatial data infrastructure (SDI)* (Phillips, Williamson et al. 1999). In this ongoing work, we propose the creation of a local, intra-organizational, service-oriented SDI as the source of detailed and specialized data on a city or a metropolitan region. This proposal extends the prevalent view on SDIs, which are mostly national or regional in scope. In any SDI, multiple information providers, each of which specializing in a set of thematic data or on data on a specific region, catalog their services in a public server, along with standard metadata. Users can then select the information services of their interest, and connect to them through the Web. This approach is beneficial in a number of ways. Firstly, users can always have access to the most up-to-date version of the data. Applications can also be kept small, with no need for large local data storage areas – an important factor for mobile computing applications. Furthermore, service-oriented architectures actually promote interoperability, since there is no need for the client application to know details on the systems that maintain the data of its interest, including data storage formats and data access methods. We also intend to contribute towards the specification of services of interest and relevance to local users, as opposed to the broader services usually considered in country-wide initiatives. This paper presents some early attempts in that direction.

This paper is organized as follows. Section 2 explores the evolution from data sharing using translation file formats, through data clearinghouses, up to SDI. Section 3 presents concepts on service-oriented software architectures. Section 4 introduces our discussion as to the functionality of a local SDI. Finally, Section 5 presents our conclusions so far and indicates a number of research directions from the themes presented here.

2. From Data Transfer Standards to Data Clearinghouses to SDI

The creation of GIS datasets is a notoriously complex and expensive undertaking. In the past, redundant efforts in dataset creation were commonplace: organizations with an interest in the same areas, therefore potential partners for sharing basic data, would not cooperate due to their diverse technological strategies, budgeting, and timing. Many potential data providers were not as keen or as quick to adopt GIS technology as their clients, so a lot of data conversion was required. This fact led to the appearance of a large data conversion market, in which private companies would convert existing maps to digital form, mostly covering situation in which basic cartography was only available on paper.

Of course, such a level of redundancy was incompatible with some budgets. As a result, in some places cooperation agreements fostered the development of large communities of users, which started to interact with each other in order to achieve common goals, such as producing common datasets, while avoiding redundant efforts. An example of such an arrangement took place in Belo Horizonte, Brazil: a broad cooperation agreement, involving 29 different organizations, including municipal, state and federal government agencies, universities, and private-sector companies has been in place for the last 13 years, and is still active (Davis Jr. and Fonseca 2005). These users interact regularly (twice monthly), have established thematic groups of interest, and have recently assembled 350 people for an two-day annual meeting, in which all projects under development were presented in 15-minute sessions.

Broad interest settings such as Belo Horizonte's rely on much effort on the part of those who coordinate data gathering and distribution. In Belo Horizonte's case, users have agreed on a dissemination structure based on a FTP server, later encapsulated under a Web page, through which common data can be downloaded in one of five different data formats. There is also a metadata sheet for each information class, even though this metadata follows no established standard and is presented as a non-searchable text file. Updating data on the FTP server requires uploading, which occurs whenever the provider of some information class finds it necessary to do so, or when one of the cooperation's partners requires it (Davis Jr. and Fonseca 2005).

Such an arrangement relies heavily on constant and efficient interpersonal communication, so that one organization's requirements on some data item can be informed to the organization that is responsible for the maintenance of that data. Therefore, even though this kind of setting is a great improvement on the possibilities that were in place in the early years of GIS usage, its potential for growing and for reaching a large number of users, including private citizens and companies, is rather limited.

2.1 Data transfer standards

A setting such as the one in Belo Horizonte depends fundamentally on data translation, since each organization can potentially use a different GIS software package. Many efforts in the past have attempted to establish a neutral file format for exchange purposes, so that every GIS would only need translators to and from this common format (Lima, Câmara et al. 2001). However, this kind of approach tends to address syntactic concerns only, avoiding semantic issues. Furthermore, such file formats are unsuitable for online access, thereby maintaining the need for an export-import offline cycle. As a consequence, multiple copies of the same data are distributed among interested parties at different times, generating serious synchronization problems.

Considering these issues, it is easy to see why none of the proposed data transfer standards managed to achieve widespread acceptance, with the possible exception of standards established by force of law or governmental regulations, such as SDTS (USGS 1998). In practice, commercial formats are used in most data transfer situations, reflecting the influence of the user community of a given GIS package. In many situations, developers tried to remedy that by distributing translators to and from the standard format for each of the most important industry-defined *de facto* standards

(Lima, Câmara et al. 2002), but that still did not solve the semantics problem nor did it point to an approach that would allow online access to various data sources.

2.2 Spatial data clearinghouses

From the establishment of a standard (or, at least, from some *de facto* standards), many national mapping agencies started to create *spatial data clearinghouses*, Internet-based components that intended to facilitate access to spatial data, by establishing a centralized site from which data from several sources can be found, and by providing complementary services, including searching, viewing, transferring, and ordering spatial data (Crompvoets, Bregt et al. 2004). Clearinghouses allow data providers to make their offerings known by users, along with descriptions of the data (metadata) and with instructions on how to access and use the data.

According to (Crompvoets, Bregt et al. 2004), there are several different understandings as to what is the definition of a spatial data clearinghouse. More recently, clearinghouses have been described in way that is very similar to the concept of a Web portal, i.e., a site in which a range of commonly used services is offered, or a gateway through which these services can be accessed (INSPIRE Architecture and Standards Working Group 2002). The emphasis on services is recent, as compared to previous definitions, which were mostly based on a combination of technical tools, institutional cooperation mechanisms, and commercial concerns (FGDC 1997).

The first spatial data clearinghouse known has been created by the United States Federal Geographic Data Committee (FGDC) in 1994, as part of an effort to combat redundant efforts on data collection. A few months later, in Brazil, the first source of publicly available spatial data on the Internet became operational: the GeoMinas project (www.geominas.mg.gov.br). GeoMinas materialized a large effort in discovering, cataloging, preparing and disseminating commonly used data on the state of Minas Gerais, Brazil. Its thematic groups discussed the creation of a shared metadata catalog, a data exchange format, the contents of a common state base map, and research on further technological means for information dissemination, including Web GIS.

Examples of national spatial data clearinghouses are the National Spatial Data Clearinghouse (USA), the GIgateway (UK), the Nationaal Clearinghouse Geo-Informatie (The Netherlands), and the Australian Spatial Data Directory (Australia). A recent Web survey analyzed and compared characteristics of 67 national spatial data clearinghouses, along with another 13 projects for implementation (Crompvoets, Bregt et al. 2004) (no Brazilian initiative was included, since, to the best of our knowledge, none exists so far). This study concluded that there is a growing dissatisfaction among clearinghouse users as to their functional capabilities, and indicated that the focus should change from a data-oriented to a user- and application-oriented view, something that can be achieved by using service-based architectures in SDI.

2.3 Spatial data infrastructures

According to (Maguire and Longley 2005), the expression “spatial data infrastructure” was proposed by the Mapping Sciences Committee of the U.S. National Research Council in 1993. It was initially used to describe the provision of standardized access to geographic information, but most of the debate on that concept reflects the ideal contents of a national SDI (or NSDI, which is actually the acronym for the American

National Spatial Data Infrastructure, “defined as the technologies, policies, and people necessary to promote sharing of geospatial data through all levels of government, the private and non-profit sectors, and the academic community”[2], and created in 1994).

Many clearinghouse initiatives evolved to what Masser (1999) calls “the first generation of national spatial data infrastructures”, while observing that the use of the term “infrastructure” implies the existence of some sort of coordination for policy formulation and implementation. Examples include the Australian SDI, the Canadian Geospatial Data Infrastructure, the Portuguese National System for Geographic Information (SNIG, Sistema Nacional de Informação Geográfica), the Malaysian National Infrastructure for Land Information Systems, and, of course, the American NSDI.

This first generation of SDI focuses on granting a broad thematic scope, which is in accordance to the objectives that allow for an analogy between SDI and other types of infrastructure: fostering economic development by means of granting access to publicly-available and multiple-use goods or services. By “publicly-available”, we do not mean “government-supported”, implying that services in an SDI can be provided by governmental agencies or by private companies, or even by private citizens, and financially supported either through an usage fee, by governmental funding, or by other economic arrangements. Even though SDIs are seen as drivers of economic development, some of the initiatives reviewed by Masser (1999) do not grant access to the private sector, or do so by charging an usage fee, as a means to recover some or all the costs of data creation and dissemination. On the other end of the spectrum is the American legal requirement to make their agencies’ data available to the public essentially for free.

Many lessons were learned from this first generation of SDIs. As Maguire et al. observe, (2005), in the USA the focus has been predominantly technical, as a result of the centralized control from the Unites States Geological Service. This resulted in a lack of attention to potential applications in governance and policy, which contributed to an unsatisfactory acceptance across governmental branches and on the private sector. Masser (1999) also observed that as a common factor over most national SDIs included in his study, and pointed out that in the future “awareness, no only within governments but also within the public at large, is likely to be the critical factor in the success of these strategies”.

The required evolution was made possible by the rapid evolution of Web-based information systems in the last few years. In the USA, the NSDI initiative went through a revision in 2002, under the federal administration’s e-government program. From that, a new initiative, called Geospatial One-Stop (GOS), was created, in order to provide widespread access to geographic information through a Web-based portal (http://www.geo-one-stop.gov). This initiative inaugurates the current concept of *geoportals* (Maguire and Longley 2005; Tait 2005).

---

[2] http://www.fgdc.gov/nsdi/nsdi.html

2.4 GeoPortals

The term "portal" has been widely used in the last few years with the general meaning of an "entry point" for information and services available on the Web, i.e., a Web site through which many other sites can be reached. Portals are, generally speaking, organized collections of references to items of interest to users. Applying this concept to geoinformation, a geoportal is, therefore, a "Web site that presents an entry point to geographic content on the Web" (Tait 2005). The intended functionality of a geoportal includes (1) the discovery of information sources and content, and (2) online access to applications.

Examples of currently existing geoportals are the previously mentioned Geospatial One-Stop[3], from the USA, the National Geospatial Data Framework (Beaumont, Longley et al. 2005) and the MultiAgency Geographic Information for the Countryside (MAGIC)[4] (Askew, Evans et al. 2005), from the UK, and the EU-Geoportal, a part of the Infrastructure for SPatial Information in Europe (INSPIRE) project (INSPIRE Architecture and Standards Working Group 2002)[5].

It is important to establish a distinction between the concepts of SDI and geoportal. We consider that an SDI is formed by the confluence of (potentially) several geographic data providers, each of which grants access to the data through specific *Web services*, software applications from which interfaces and bindings are expressed in XML and that can be discovered using XML messages (W3C Web Services Architecture Working Group 2002). In order to select which data, and, as a consequence, which services should be accessed to fulfill his needs, the user or client searches through a repository of metadata on available geographic data and services. Naturally, the providers of such data and services need to have the corresponding metadata previously included in this repository. In the case of a human user, searches are done interactively, through a geoportal, possibly using search interfaces and other interactive tools; in the case of a software client, this can be done through a catalog Web service. With this, we consider that a geoportal should be considered to be a component of an SDI (Figure 1).

---

[3] http://www.geo-one-stop.gov

[4] http://www.magic.gov.uk

[5] http://eu-geoportal.jrc.it

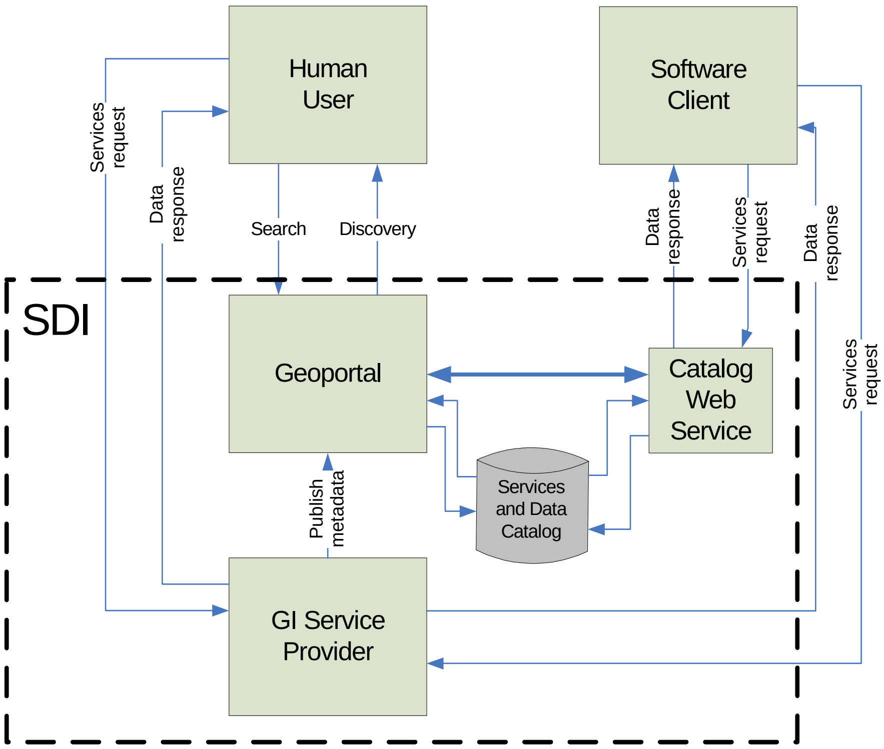

Figure 1 - Geoportals and SDI

The use of Web services to grant direct access to data is the most important distinction between first- and second-generation SDIs. In fact, the numerous possibilities that arise from using services to encapsulate data from multiple sources, and thereby achieve interoperability, have led Bernard and Craglia (2005) to propose a new translation for the SDI acronym: *Service-Driven Infrastructures*.

In this work, our intention is to explore the concept of SDI further, in order to be able to propose services that are specific to local users and applications. We intend to do so by proposing basic services and by designing ways to implementing the composition of them into more complex services. But before we go into that, we must first present concepts on service-oriented architectures and Web services, including the Open Geospatial Consortium's view on them (Percivall 2003). In fact, a service-based architecture is being use in OGC's Geospatial Semantic Web Interoperability Experiment, launched in April 2005 (Lieberman, Pehle et al. 2005).

3. Service-based distributed system architectures

Component-based software development is not recent, and has been the subject of much interest nowadays because of its potential to reduce development costs and time, and because of the interest in the deployment of distributed systems. One of the most interesting approaches in this field is the one of *service-oriented architectures* (SOA) (Papazoglou and Georgakopoulos 2003).

Services, along with their descriptions and fundamental operations, such as *discovery*, *selection* and *binding*, constitute the basis of SOA. SOA supports large applications with sharing of data and processing capacity, through network-based distributed allocation of applications and use of computational resources. In this architecture, services are self-contained, which means that information on the service's description, including its capabilities, interface, behavior, and quality, can be obtained from the service itself, through a standardized set of functions. Services can either be *primary*, i.e., developed using some programming language, or they can be formed by the composition of other services (Curbera, Khalaf et al. 2003).

Service providers, service aggregators and service users are the actors that participate in this scenario. Providers implement and publish services, while aggregators design compositions of rules based on primary services. The available services should be listed in directories for user reference, or "discovery". Service users may be human or software clients, which need to access the services through the communications network. The sharing of computational resources that is provided by this architecture allows the use of "thin clients", devices that usually have restricted storage and processing capacity.

3.1 Web services

Web services are a particular class of services that use open Internet standards, such as connection and communication using the Hypertext Transfer Protocol (HTTP), identification using the Uniform Resource Identifier (URI), contents specification through the eXtensible Markup Language (XML), service descriptions expressed by the Web Services Definition Language (WSDL), and directory services using the Universal Description, Discovery and Integration (UDDI) protocol (W3C Web Services Architecture Working Group 2002; Ferris and Farrell 2003).

Therefore, while services in general provide interoperability between different software components, Web services go a step further by facilitating cross-institutional interchange of data and services over the Internet, and by improving the sharing of resources among a variety of data sources.

The main Web service operations are *publication*, *discovery*, *selection*, *binding*, and service *composition*. *Publication* is performed by a service provider, and consists in the creation of a service description in WSDL and its publishing on discovery channels on the Web. These channels use the UDDI protocol to register the service, along with some details that allow users to *discover* it. Users can then *select* services using UDDI to search through catalogs or directories. Search results are actually sets of URIs. Once the service descriptor is obtained (in WSDL) for the *binding* operation, client software initiates a direct communication with the service provider through the Internet using HTTP, sending a request. The binding ends with the reception of the expected Web service response in XML (Figure 2). Finally, a set of more complex services can be created by the *composition* (or *chaining*) of primary ones. Compositions are designed mainly by service aggregators. In this case, service users access the composite services in the same way as simple services, but different results are possible. Thus, several alternatives in terms of service performance, costs or quality become feasible through the appropriate configuration of Web service chains.

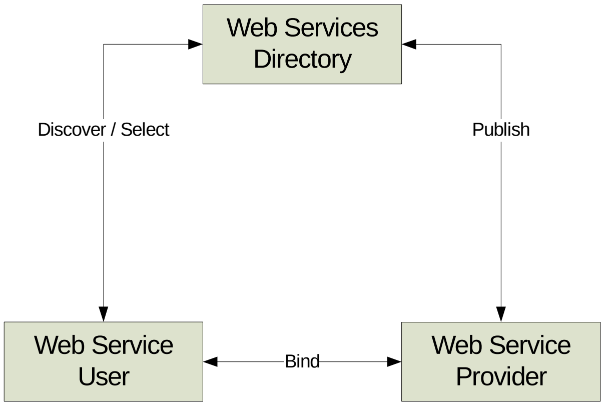

Figure 2 - Web service operations

3.2 OGC Web Services

The Open Geospatial Consortium (OGC) proposed an architecture for sharing of geographic data and functionality over the Internet, thus leading the standardization process regarding data formats, methods and interface specifications. This architecture is called the *OpenGIS Services Framework* (Percivall 2003).

The OpenGIS Services Framework does not necessarily use the usual Web services standards, such as the Simple Object Access Protocol (SOAP) and WSDL. Nevertheless, the adoption of these standards is desirable to ensure interoperability between regular Web services and OGC Web services. Instead of using UDDI, the OGC proposes the use of catalog services for the implementation of the publication, discovery and selection operations. Moreover, OGC Web services have a particular interface for binding, which does not use service descriptors. This alternative poses difficulties for the indexing and searching of services. Also, OGC Web services use Geographic Markup Language (GML) to encode and transmit objects, while regular Web services use generic XML, but this is not actually a difference, since GML is based on XML. We observe that the main differences that exist between W3C Web services and the OGC ones should be solved as soon as possible, in order to incentive the adoption of the proposed standards in a more universal way (Sonnet 2004).

Some basic Web services were specified by OGC, as services applied to registry, composition, visualization and codification. The main Web services are described below (Davis Jr, Borges et al. 2005).

- **Web Feature Service:** provides a interface for the insertion, selection, updating and removal of geographic features (objects).
- **Web Coverage Service:** provides access to geo-fields, much in the same manner of the Web Feature Service. Notice that this service does not return images of the geo-fields, but rather returns semantic details on them.
- **Web Gazetteer Service:** extends the Web Feature Service with resources for the implementation of interfaces to gazetteers (Souza, Davis Jr et al. 2005). This service is still under discussion at the OGC.
- **Web Registry Service and OpenGIS Catalog Service (OCS):** implement an operational functionality similar to UDDI.

- **Web Coordinate Transformation Service:** provides an algorithm that converts coordinates for spatial objects between different spatial reference systems.
- **Web Map Service:** a service for the production of maps over the Web, or Web Maps. Maps, in this service, are renderings (presentations) of the reality, and do not include the actual geographic data.
- **Web Terrain Service:** similar to the Web Map Service, but geared towards three-dimensional renderings of surfaces. Both this and the Web Map Service can produce renderings based on image formats or on the vector-based Scalable Vector Graphics (SVG) format.

OGC Web services may also be combined. A special kind of OGC Web service, the Web Notification Service, can be used to send update notifications to registered clients that participate on a chain in which some Web service is to be altered. Interoperability interfaces are also being developed by the OGC for sensor networks, under names such as Sensor Collection Service and Sensor Planning Service. This shows the degree of commitment of the OGC with service-oriented architectures for interoperability purposes, which historically represent the core of the OGC's purposes.

Such an open and flexible architecture will find its main uses in situations very similar to the ones presented in the discussion on SDIs. Summing up our previous arguments, SDIs must be distributed, must support multiple applications, multiple clients of several different types, multiple data sources, multiple data maintenance teams, under a heterogeneous computational environment. SDIs must not force the adoption of specific *products* on their participants, but should instead provide an architectural view and determine a set of minimal standards. These standards should be as widely accepted as possible, and typical Internet standards as the ones we mentioned in this section fill that description closely.

Based on a service-oriented architecture, considering the needs of spatial data infrastructures, we now proceed to propose a set of services that must be available for local SDIs, considering the peculiar characteristics of urban spaces and the probable demands of potential urban SDI users.

4. The Design of a Local SDI

As previously mentioned, most SDIs reflect the needs of national mapping agencies (Masser 1999; Crompvoets, Bregt et al. 2004), which means dealing with large volumes of data on a relatively limited number of themes, mostly in regional scales.

The richness of local GIS applications, however, indicates that there is still much to be done in order to establish the SDI needs of local GI users. Furthermore, the diversity of possible interests of individuals as to geographic information for everyday tasks shows that a clearly established, service-based, local SDI can be the ideal starting point for initiatives such as location-based services, route planning, convenience shopping and many others. On the other hand, such a wide array of information cannot possibly be provided by the local government alone – many of the required items are the subject of interest of other actors, such as utility companies, state and federal agencies, class associations, NGOs, private companies and others (Figure 3).

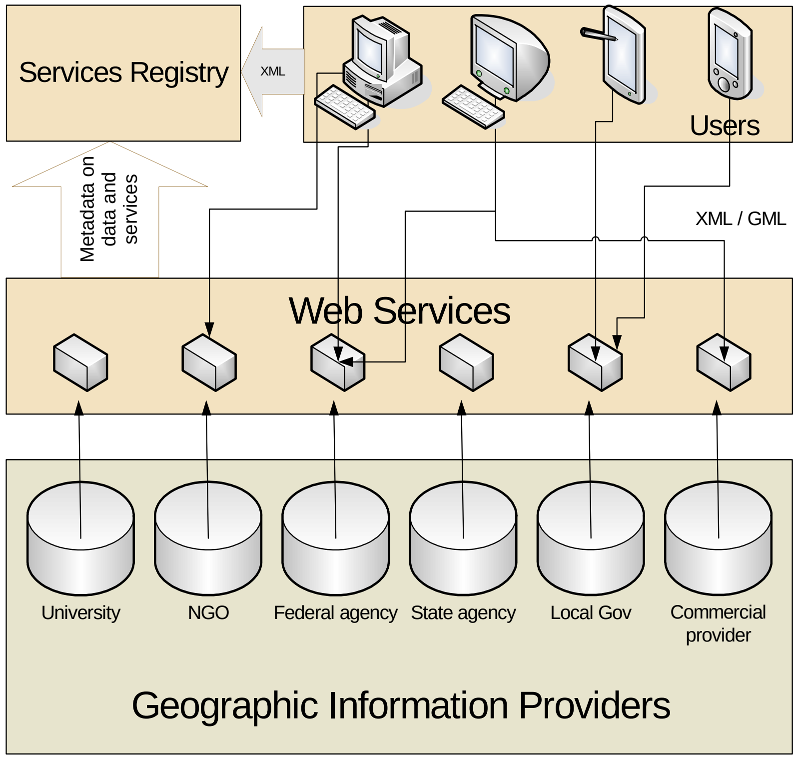

Figure 3 - A Local SDI

We also observe that the service-based implementation alternative for local SDIs has the advantage of clearly establishing a differentiation between an SDI and a local GIS. In our opinion, the concept of SDI reflects widely-available, general-use data, while a GIS is usually built in an organization around a definite goal or set of goals. As this set of goals gets more comprehensive, so do the data and the applications, thereby causing the confusion. (Jacoby, Smith et al. 2002), for instance, states that "(…) many cities or Local Governments world-wide have established their own SDIs although they are more commonly referred to as geographical information systems". We consider that a local government GIS may or may not be the basis for a local SDI, depending on how (how well, how freely) the government distributes (or not) its information. If the GIS is focused only on administrative issues, it may not fill the requirements of a true SDI, even though its informational contents may be quite comprehensive.

Thus, our discussion starts in the definition of the requirements for a local SDI, which needs to be done based on a set of fundamental services that need to be provided. Initially, it is clear to us that these services should include, but are not limited to, the following:

- **Basemap:** a service to provide access to a basic city map, with street names, neighborhoods, and so on, preferably displayed over high-resolution imagery, High-resolution images are very intuitively interpreted and recognized by the laymen, so they play a fundamental role in the popularization of spatial information (consider, for instance, Google Earth's potential impact). Notice that the client of such a service is the responsible for the displaying of map data, so that visualization parameters can be decided by each user/application. Using such a visualization tool, most of the other services can also produce viewable results, even though there can be several applications for which map viewing is not a requirement. A simple client for this service could use a free and simple SVG-based viewer (Guo, Zhou et al. 2003; Mathiak, Kupfer et al. 2004).

- **Personal location:** a service which would receive any kind of information on the location of a mobile client (nearby address, names of crossing streets, place name, or cellular base station identification) and would return a coordinate that can be used by the application in a more precise manner (see "emergency services", below). We also call that possibility a "virtual GPS". Positions returned by this service may have a certainty indicator associated with them, showing to the user how precise the results actually are (Davis Jr, Fonseca et al. 2003).
- **Geocoding/address recognition and location:** this should be one of the prominent services in an urban SDI, since addresses are the most important form of spatial location in urban areas. A geocoding service could exceed the norm by not only providing a coordinate that corresponds to a given textual address, but by being able to locate places by name as well. In this point, a geocoding service could extend the notion and incorporate the functions of a local gazetteer (Davis Jr, Fonseca et al. 2003; Souza 2005; Souza, Davis Jr et al. 2005).
- **Basic routing services:** from an origin point, expressed as an address or as a place name, this service would be able to generate a route description, a route map, or both. The route description can be provided as a walking itinerary, as a driving itinerary, or a trip using the public transportation system (see below).
- **Public transportation system:** this service would provide directions (see basic routing services) to the nearest terminals or access means to all kinds of public transportation services, including buses, subway/metro rail, taxis, and even car rental services.
- **Public services:** a service dedicated to the location of public offices, health services, post offices, utility services, and so on.
- **Private services:** a service dedicated to the location of commercial activities such as shops, restaurants, hotels, and others. In conjunction with other services, this could lead to a "visitor information" service.
- **Emergency:** a service dedicated to present *local* details on what to do in emergency situations, including the determination of the closest health facility and the route to it.

These proposed services differ from the ones that are the target of NSDIs, since they require a greater level of detail as to basic data, and the potentially require access to several different data sources, each of which maintained by a different data provider. Also notice that many of those services would be ideally suited for mobile computing applications – the progress of which has been severely hindered by the fact that the telecom companies, which hold the customer base and the communications network, do not usually have access to constantly updated quality data, and have little interest in developing such systems by themselves. An SDI equipped with such information services will bring us a significant step closer to Max Egenhofer's proposals for spatial information appliances (Egenhofer 1999).

Another range of services that could be incorporated to local SDIs is the public access to government-owned spatial information, in fields such as urban cadastre and property registration, public works, and environmental control. The first corresponds to a subject in which the interface between local governments and private citizens is most intensive. As of the latter, we observe that citizen participation is an asset to achieving significant results in environmental protection, but that participation can be much enhanced by the use of information technology tools.

Also notice that this proposal can be developed based on the OGC Web Services architecture, but detailing and adjusting the services for typical urban uses, to be based on sufficiently detailed data. The interconnection and chaining between services is something that is also interesting for the research (Einspanier, Lutz et al. 2003; Granell, Gould et al. 2005), both because it facilitates the development of more complex services over simpler ones, and because it can enable the development of regional and national SDIs over local services (Rajabifard, Williamson et al. 2000; Rajabifard and Williamson 2001), possibly using techniques from the field of multiple representations in GIS (Davis Jr and Laender 1999). Current limitations of OGC Web Services in that regard are being addressed in the OGC Web Services Phases 2 (OWS-2) and 3 (OWS-3), which intend to extend the existing standards and to make them compatible with the World Wide Web Consortium's usual standards, including SOAP and WSDL (Lieberman, Pehle et al. 2005).

5. Conclusions and future work

This is a work in progress. We are now at the early stages of designing a prototype SDI, in order to be able to implement and test many of the ideas and concepts presented here as to their feasibility, performance, consumption of computational resources, and ease of dissemination. From there, we will implement basic Web services on address location, geocoding, and routing, from which more complex services and applications can be designed. We intend to (1) develop a better understanding of the SDI approach, including its demands for computational resources, (2) develop a method for the design and development of Web services connected to the SDI, (3) design means to publish metadata on services, and (4) study the architectural possibilities for chained Web services as applied to SDI, with a focus on distributed data management and distributed processing. The latter of these objectives should provide an interesting framework for mobile GIS applications, a field in which local (urban) geographic knowledge is indispensable. In this research, important issues, such as privacy and security, will be taken into consideration. Our testbed will include data and services from several different organizations, participants of Belo Horizonte's cooperation agreement for information exchange.

References

Askew, D., Evans, S., Matthews, R. and Swanton, P. (2005). "MAGIC: a geoportal for the English countryside." Computers, Environment and Urban Systems **29**(1): 71-85.

Beaumont, P., Longley, P. A. and Maguire, D. J. (2005). "Geographic information portals - a UK perspective." Computers, Environment and Urban Systems **29**(1): 49-69.

Bernard, L. and Craglia, M. (2005). SDI - From Spatial Data Infrastructure to Service Driven Infrastructure. Research Workshop on Cross-Learning Between Spatial Data Infrastructures and Information Infrastructures, Enschede, The Netherlands.

Crompvoets, J., Bregt, A., Rajabifard, A. and Williamson, I. (2004). "Assessing the worldwide developments of national spatial data clearinghouses." International Journal of Geographic Information Science **18**(7): 665-689.

Curbera, F., Khalaf, R., Mukhi, N., Tai, S. and Weerawarana, S. (2003). "The next step in Web services." Communications of the ACM **46**(10): 29-34.

Davis Jr, C. A., Borges, K. A. V., Souza, L. A. d., Casanova, M. A. and Lima, P. (2005). O Open Geospatial Consortium. Bancos de Dados Geográficos. M. A. Casanova, G. Câmara, C. A. Davis Jr, L. Vinhas and G. R. d. Queiroz. Curitiba (PR), Editora MundoGEO**:** 379-395.

Davis Jr, C. A., Fonseca, F. T. and Borges, K. A. V. (2003). A flexible addressing system for approximate urban geocoding. V Brazilian Symposium on GeoInformatics (GeoInfo 2003), Campos do Jordão (SP).

Davis Jr, C. A. and Laender, A. H. F. (1999). Multiple Representations in GIS: materialization through map generalization, geometric, and spatial analysis operators. 7th International Symposium on Advances in Geographic Information Systems (ACM GIS), Kansas City, MO.

Davis Jr., C. A. and Fonseca, F. T. (2005). "Considerations from the Development of a Local Spatial Data Infrastructure." Information Technology for Development(to appear).

Egenhofer, M. (1999). Spatial information appliances: a next generation of geographic information systems. I Brazilian Symposium on GeoInformatics (GeoInfo 1999), Campinas (SP).

Einspanier, U., Lutz, M., Senkler, K., Simonis, I. and Sliwinski, A. (2003). Towards a Process Model for GI Service Composition. GI-Tage 2003, Münster, Germany.

Ferris, C. and Farrell, J. (2003). "What Are Web Services?" Communications of the ACM **46**(6): 31.

FGDC (1997). Metadata to Clearinghouse Hands-On Tutorial. Washington, DC, Federal Geographic Data Committee.

FGDC (2001). Content Standard for Digital Geospatial Metadata Workbook. Reston, VA, Federal Geographic Data Committee.

Fonseca, F., Egenhofer, M. and Borges, K. A. V. (2000). Ontologias e Interoperabilidade Semântica entre SIGs. II Workshop Brasileiro de GeoInformática (GeoInfo 2000), São Paulo (SP).

Granell, C., Gould, M. and Ramos, F. (2005). Service Composition for SDIs: integrated components creation. 2nd International Workshop on Geographic Information Management (GIM'05), Copenhagen, Denmark.

Guo, Z., Zhou, S., Xu, Z. and Zhou, A. (2003). G2ST: A novel method to transform GML to SVG. 11th ACM International Symposium on Advances in Geographic Information Systems (ACM GIS), New Orleans (LA).

INSPIRE Architecture and Standards Working Group (2002). INSPIRE Architecture and Standards Position Paper. Brussels, Commission of the European Communities.

Jacoby, S., Smith, J., Ting, L. and Williamson, I. (2002). "Developing a common spatial data infrastructure between State and Local Government -- an Australian case study." International Journal of Geographic Information Science **16**(4): 305-322.

Lieberman, J., Pehle, T. and Dean, M. (2005). Semantic Evolution of Geospatial Web Services. W3C Workshop on Frameworks for Semantic Web Services, Innsbruck, Austria, The World Wide Web Consortium (W3C).

Lima, P., Câmara, G. and Monteiro, A. M. V. (2001). Intercâmbio de Dados Geográficos: Modelos, Formatos e Conversores. III Workshop Brasileiro de GeoInformatica (GeoInfo 2001), Rio de Janeiro (RJ).

Lima, P., Câmara, G. and Queiroz, G. R. d. (2002). GeoBR: Intercâmbio Sintático e Semântico de Dados Espaciais. IV Brazilian Symposium on GeoInformatics (GeoInfo 2002), Caxambu (MG).

Maguire, D. J. and Longley, P. A. (2005). "The emergence of geoportals and their role in spatial data infrastructures." Computers, Environment and Urban Systems **29**(1): 3.

Masser, I. (1999). "All shapes and sizes: the first generation of national spatial data infrastructures." International Journal of Geographic Information Science **13**(1): 67-84.

Mathiak, B., Kupfer, A. and Neumann, K. (2004). Using XML languages for modeling and Web-visualization of geographical legacy data. VI Brazilian Symposium on GeoInformatics (GeoInfo 2004), Campos do Jordão (SP).

Papazoglou, M. P. and Georgakopoulos, D. (2003). "Service-Oriented Computing." Communications of the ACM **46**(10): 25-28.

Percivall, G. e. (2003). OpenGIS Reference Model, Open Geospatial Consortium, Inc.

Phillips, A., Williamson, I. and Ezigbalike, C. (1999). "Spatial Data Infrastructure concepts." The Australian Surveyor **44**(1): 20-28.

Rajabifard, A. and Williamson, I. P. (2001). Spatial Data Infrastructures: Concept, Hierarchy, and Future Directions. GEOMATICS'80, Tehran, Iran.

Rajabifard, A., Williamson, I. P., Holland, P. and Johnstone, G. (2000). From Local to Global SDI Initiatives: a pyramid of building blocks. 4th Global Spatial Data Infrastructure Conference, Cape Town, South Africa.

Sonnet, J. (2004). OWS 2 Common Architecture: WSDL SOAP UDDI., Open Geospatial Consortium, Inc.

Souza, L. A. (2005). Locus: Um Sistema de Localização Geográfica através de Referências Espaciais Indiretas. Departamento de Ciência da Computação. Belo Horizonte (MG), Universidade Federal de Minas Gerais.

Souza, L. A., Davis Jr, C. A., Borges, K. A. V., Delboni, T. M. and Laender, A. H. F. (2005). The Role of Gazetteers in Geographic Knowledge Discovery on the Web. 3rd Latin American Web Congress (LA Web 2005), Buenos Aires, Argentina.

Tait, M. G. (2005). "Implementing geoportals: applications of distributed GIS." Computers, Environment and Urban Systems **29**(1): 33-47.

USGS (1998). Spatial Data Transfer Standard. Rolla, MO, United States Geological Survey.

W3C Web Services Architecture Working Group (2002). Web Services Architecture Requirements, W3C Working Draft, World-Wide Web Consortium.