Service Chaining of Interoperable Geographic Information Web Services

By
Nadine Alameh
Global Science & Technology
6411 Ivy Lane Suite 300, Greenbelt, MD 20770
alameh@gst.com

Keywords: Web services, service chaining, distributed GIS, GIS evolution

Service Chaining of Interoperable Geographic Information Web Services

This paper is based on the evolution of Geographic Information Systems towards the web services model. Indeed, this model is rapidly materializing as a result of the advancements in general web service technologies and the focused efforts of the Open GIS Consortium in the areas of service categorization and interoperability of service interfaces. This paper focuses on the issues of service chaining, the process of combining or pipelining results from several complementary and interoperable GIS Web Services to create a customized solution. The paper presents three basic service chaining options: client-coordinated, static and mediated chaining. Each option is examined with respect to its coordination complexity, metadata tracking ability and error handling propagation. The paper highlights the middleware requirements specific to working with GIS services. The paper concludes with a discussion on the potential changes in the GIS marketplace given a successful deployment and chaining of GIS Web Services.

1 Introduction

Over the last decade, we have witnessed the evolution of Geographic Information Systems (GIS) from the traditional model of stand-alone systems with geo-data tightly coupled with the systems used to create them, to an increasingly distributed model based on independently-provided, specialized interoperable GIS Web Services (Alameh 2001). This evolution is enabled by the advancements in supporting IT technologies and the growing demand for GIS in a variety of application domains (Figure 1).

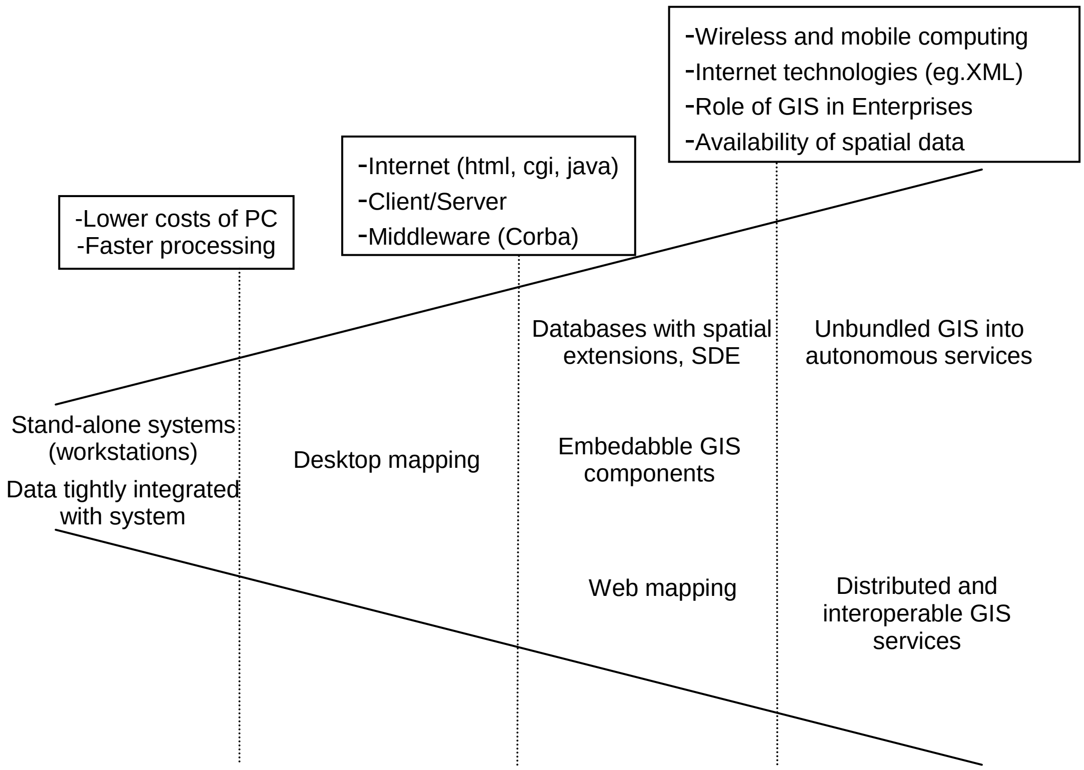

Figure 1. Evolution of GIS.

This gradual on-going transformation is primarily fueled by the growing role of GIS in today's organizations, the increasing availability of spatial data and its inherent conduciveness to reuse, the relative maturity of web and distributed computing technologies, and the key role of GIS in a promising location-based services market (Figure 2).

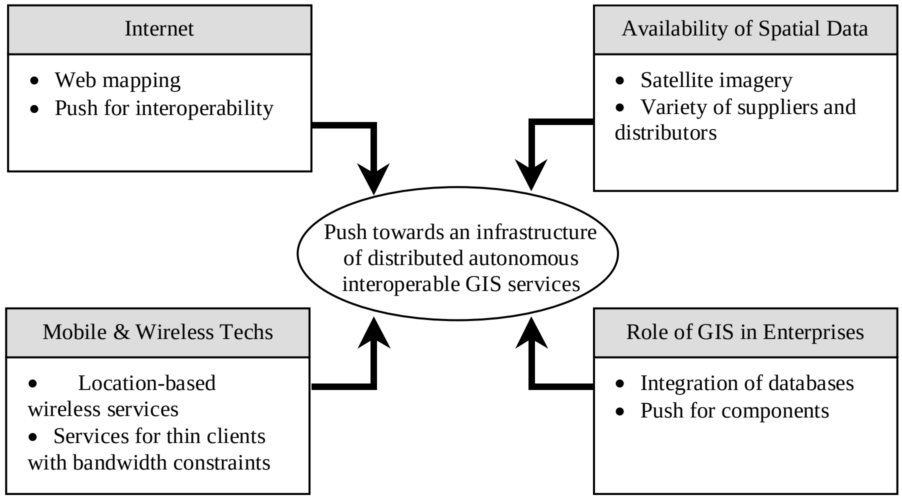

Figure 2. Key drivers of GIS Web Services Architecture.

The service-based model is rapidly materializing as a result of the advancements in general web service technologies and the focused efforts of the Open GIS Consortium (OGC) in the areas of service categorization and interoperability of service interfaces (Buehler and McKee 1998). Soon, it will be possible to dynamically assemble applications from multiple network-enabled GIS services for use in a variety of client applications. This dynamic assembly of services is the motivation behind studying service chaining in this paper. Service chaining is the process of combining or pipelining results from several complementary GIS Web Services to create customized applications.

After an overview of the GIS Web Services architecture, the paper presents three basic service chaining options: client-coordinated chaining, static chaining and mediated chaining. Each option is examined with respect to its coordination complexity, metadata tracking ability and error handling propagation. The paper highlights the middleware requirements specific to working with GIS services. The paper concludes with a discussion on the potential changes in the GIS marketplace given successful deployment and chaining of GIS Web Services.

2 Overview of GIS Web Services Architecture

A simplified view of the GIS Web Services Architecture is presented in Figure 3. This figure shows a variety of GIS Web Services (defined in Section 2.1) being chained and accessed via standardized interfaces by a range of clients (discussed in Section 2.2).

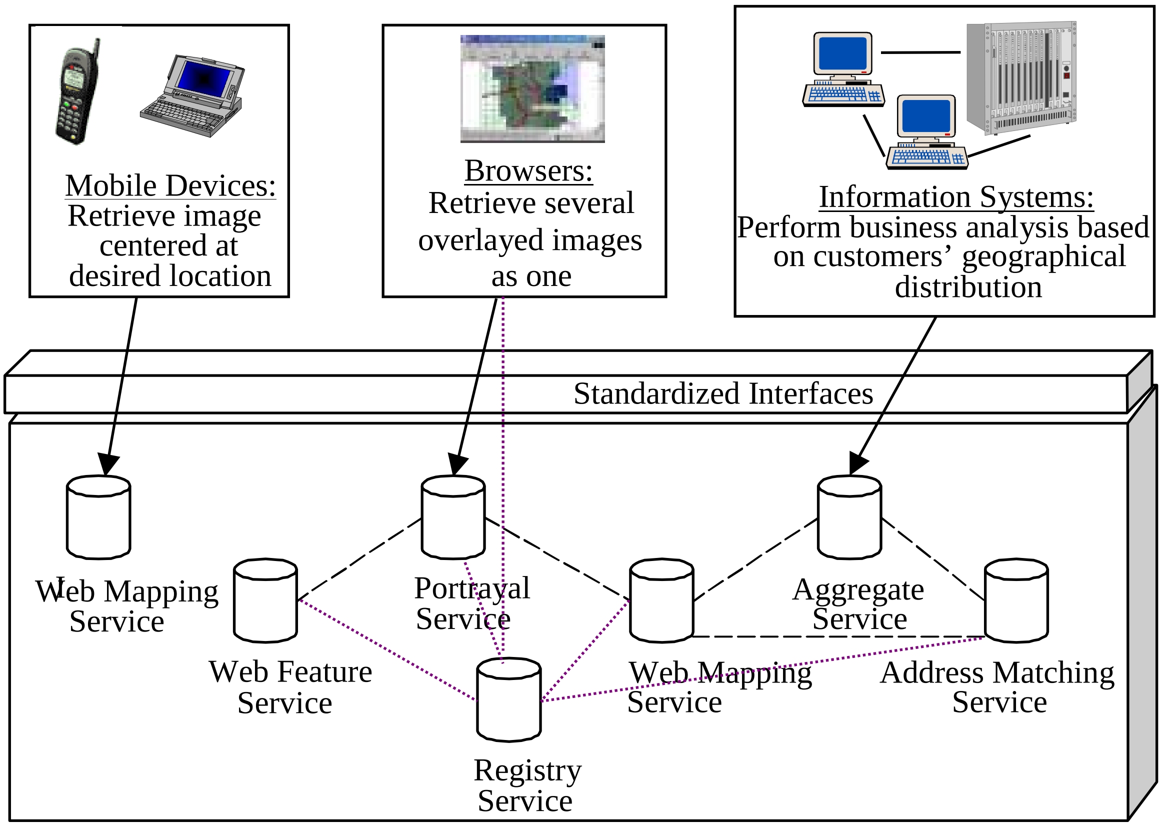

Figure 3. Simplified View of GIS Web Services Architecture.

2.1 *GIS Web Services*

Web services are self-contained, self-describing, modular applications that can be published, located, and dynamically invoked across the web (ISO 2001). Web services provide access to sets of operations accessible through one or more standardized interfaces. In the process, services may use other external services and operations. Services can be grouped into three basic categories:

- **Data services** (such as the OGC Web Mapping, Web Coverage and Web Feature services) offer customized data to users (OGC 2002). These services are tightly coupled with specific data sets.
- **Processing services** are not associated with specific datasets. Instead they provide operations for processing or transforming data in a manner determined by user-specified parameters (Alameh 2001). Such services can provide generic processing functions such as projection/coordinate conversion, rasterization/vectorization, map overlay, imagery manipulation, or feature detection and imagery classification.
- **Registry/catalog services** are used to classify, register, describe, search, maintain and access information about Web services (OGC 2002). Types of registries are differentiated by their role such as registries for cataloguing data types, online data instances, service types and online service instances.

In order for a sustainable and extensible GIS Web Services Architecture to exist, the basic services should be accessed via standardized interfaces, which are currently being developed within OGC and are not in the scope of this paper.

2.2 *Client Applications*

Once GIS Web Services are deployed, client applications can be built more flexibly by mixing and matching available services. Figure 3 shows a variety of clients that can be directly linked to the key distribution drivers presented in Figure 2. The clients featured cover a range of possible applications: thin clients such as web browsers, newly emerging clients such as hand-held devices as well as the traditional information and GIS systems.

Each client application is created by using multiple GIS and non-GIS Web Services. Service chaining is an integral part of developing customized client applications in such an environment. Figure 4 shows a simple service chaining example where several GIS coverages are fetched from Web Coverage Services (WCS), mosaicked and portrayed using a Coverage Portrayal Service (CPS). The resultant coverage is then re-projected to another Spatial Reference System (SRS) using a processing service. The coverage along with feature data extracted from a Web Feature Service (WFS) are then overlayed and portrayed as a finished image to the client. For more information on WCS, CPS and WFS, refer to the OGC webpage (OGC 2002).

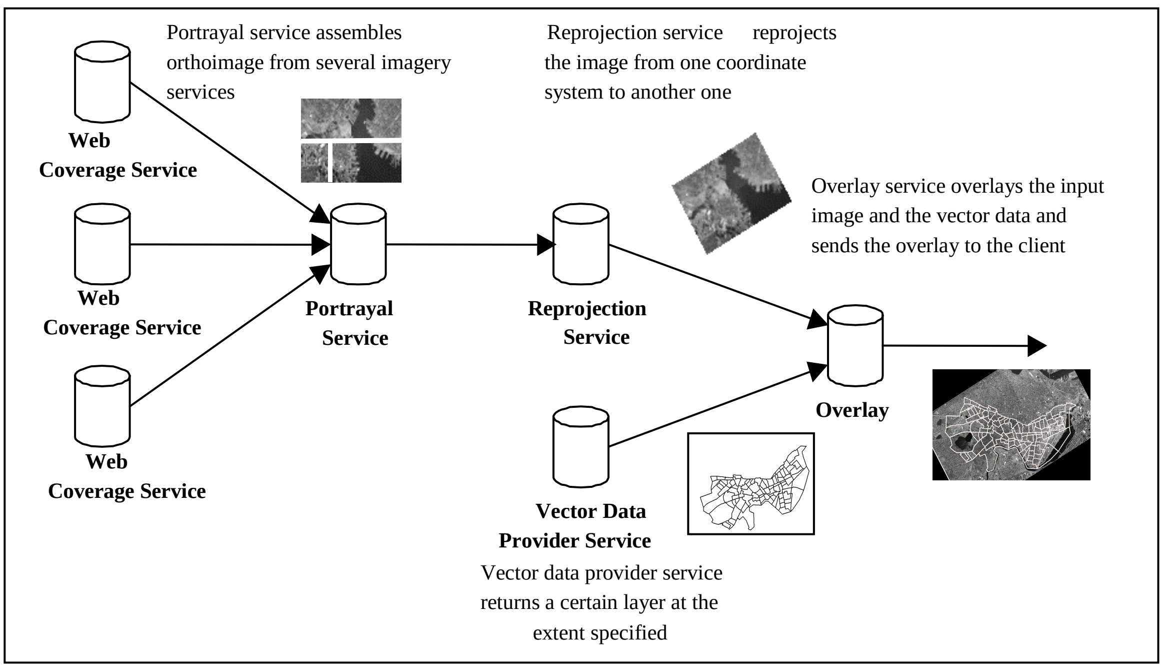

Figure 4. A Simple Service Chaining Example.

There are several alternatives for efficient and practical dialogs that occur among the individual services being chained as well as between the services and the calling client. Using a simple example, the next section presents three basic alternatives for GIS Web Services chaining.

3 Service Chaining

3.1.1 Example

To establish a basis for discussing the different alternatives, we construct a simple yet non-trivial example, and examine it under the light of each service chaining option. We use the example of a user providing an address to an application and requesting a geo-referenced image centered at that address (for the sake of simplicity, we assume that the size of the image in pixels is fixed). One aspect of this example's simplicity is that the GIS data types handled are limited to raster imagery. Therefore, the example avoids the additional complexities of heterogeneous semantics and topology representations. Nevertheless, despite its simplicity, the example is rich enough for exploring the trade-offs in various service chaining approaches.

The service types needed in the example are:

- **An address matching service** (e.g. the Etak service at www.etak.com): according to a GeocoderService RFC draft submitted to OGC, an address matching service transforms a phrase or term that uniquely identifies a feature such as a place or an address into applicable geometry (usually either a coordinate x,y or a minimum bounding rectangle). For simplicity, we assume that the service used by our client provides (x,y) coordinates in any spatial reference system. Given that typical address matching services also return additional information such as normalized address, matching precision and location precision, we assume that this additional information can be filtered out such that the client receives only the coordinates in response to a request. In some cases, address matching services return several locations matching a given address. In such cases, user intervention may be needed to identify the intended address.

- **A Web Mapping service**: returns a map corresponding to pre-specified geographic and pixel dimensions of an area (Doyle 2000).
- **A re-projection service**: transforms a raster image from one spatial reference system to another. This service is needed because the native projection of a data set may be different from the one requested by the client. A note here about the fact that there is no re-projection service standard interface yet but that we assume that they will consist of at least the following parameters.

Figure 5 provides an input/output illustration of services from the three types above. We emphasize that the client in our example is not limited to a particular service instance from each type. The client is expected to use registry/catalog services to find service instances belonging to desired service types.

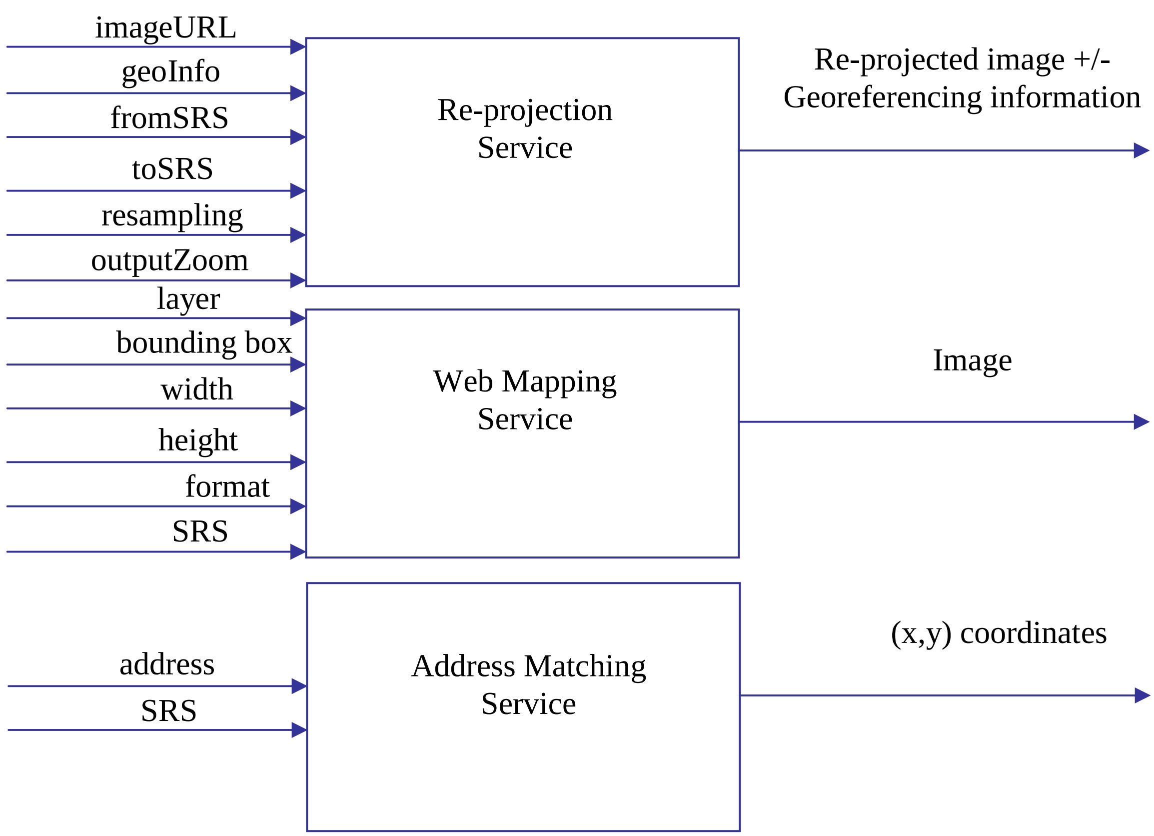

Figure 5. Illustration of Services Used in the Example (With Only Key Parameters Shown).

For the sake of simplicity, we assume that the authentication of the client by the services can be handled using available authentication technologies, such as Kerberos, cookies or basic http authentication. Similarly, we assume billing for services and data can be managed using current e-commerce related approaches, such as Ecash (www.ecash.com), CyberCash (www.cybercash.com) or PayPal (www.paypal.com).

3.2 *Key Analysis Issues*

There is a wide range of issues surrounding the design of a scalable and extensible geospatial infrastructure based on service chaining. The key is to address the issues of service chaining and coordination without compromising performance and scalability of the infrastructure. In this section, we focus on a subset of these issues. This subset covers

- The degree of transparency to the client of the service chaining complexities and the amount of coordination responsibility required by the client to construct, execute and manage service chains.
- The alternatives for keeping track and relaying to the client the sources of data used in a service chain and the transformations that were applied to them along the chain. Keeping track of metadata may be especially important in cases where service chaining is hidden from the client.
- The alternatives for handling and reporting errors along a chain to the client. This may be regarded as a special case of metadata tracking. Precise error reporting is needed by clients in the cases when one or more services along a chain break or issue a certain type of exception.

The service chaining described in Section 3.1 is used to pinpoint issues of complexity of dialogs and the implications of service chaining on the client's thickness and intelligence requirements.

3.2.1 Client-Coordinated Service Chaining

In the simplest case, service chaining is fully transparent to the client: The client defines and controls the order of execution of individual services in a chain. The client also must have prior knowledge of the interfaces, inputs and outputs of the services or service types it is using. In the upper part of Figure 6, the client first searches a Catalog to find service instances to use (Nebert 1999). A service instance for each service type needed is returned. The client then uses the Address Matching Service to find the coordinates of the address requested by the user. The coordinates (along with other parameters) are sent to a Web Mapping Service which returns an image to the client. If the image returned is in a different SRS than the one used by the client, it is then re-projected by a Re-projection Service (the client may access the registry service again to find a Re-projection Service instance). In this case, the client is handling all intermediate results returned by the services in the chain.

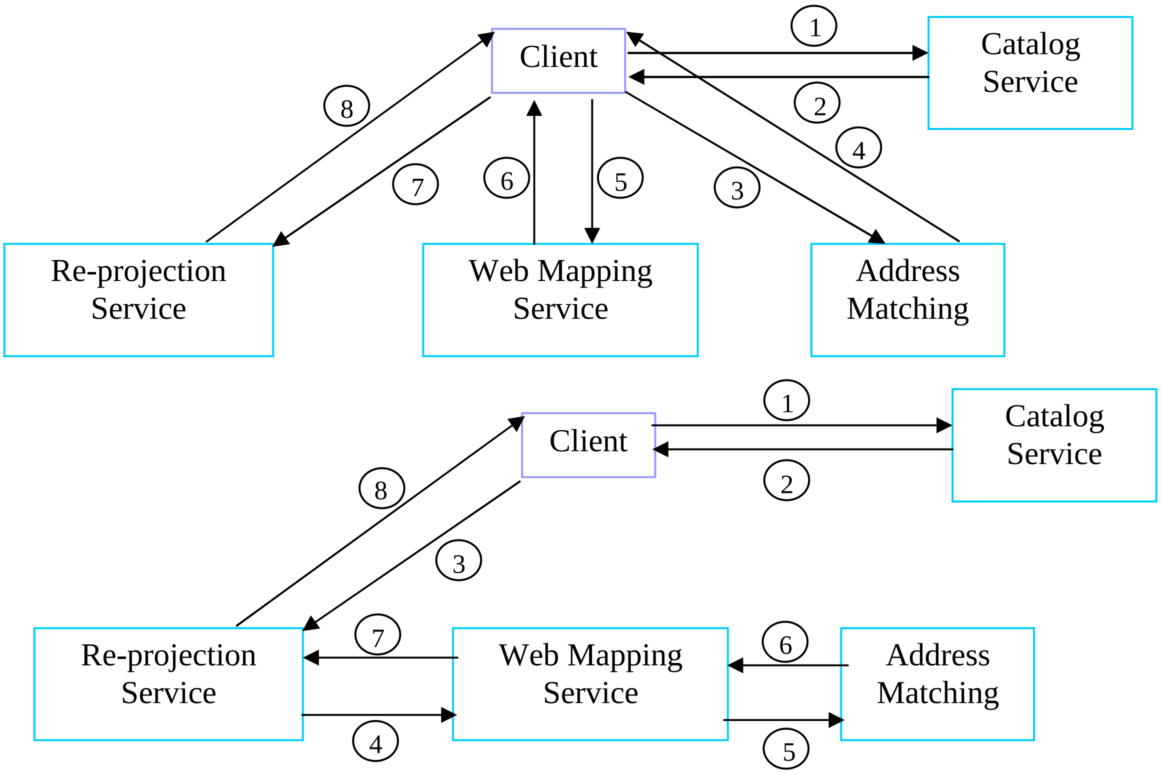

Figure 6. Client-Coordinated Service Chaining.

The lower part of Figure 6 shows how the client can directly embed, in the input to one service, a request to the following service. Instead of downloading the image itself, the client provides the Re-projection Service with the URL of the Web Mapping Service needed to retrieve the image. In turn, the bounding box sent to the Web Mapping Service is specified by the client as a function of the result of a request to the Address Matching Service. This approach of nesting calls within a service chain is consistent with the one currently followed by OGC in its initial Coverage Portrayal Service(CPS)/Styled Layer Descriptor(SLD) chains (OWS1, 2002).

While nesting calls promises to simplify some of the coordination responsibilities of the client, it nevertheless introduces complexities in the areas of error and metadata propagation as well as the client's ability to control certain details. Consider for example the case of the Address Matching service issuing an exception in response to an erroneous call from the client. To the Web Mapping Service in the lower part of Figure 6, this exception is viewed as an invalid input, triggering the Web Mapping Service to signal an "invalid input" exception. A domino effect ensues, as the same problem occurs over the Web Mapping/re-projection link, forcing the Re-projection Service to signal its own "invalid input" exception (this time to the client). Although the client is eventually informed of the occurrence of an exception in the chain, the actual exception received by the client does not disclose information about the source or the cause of that exception. One way to overcome this limitation is to allow a service to automatically forward a received exception input, as is, to the next service in the chain; while appending to the forwarded exception any of the service's own error messages. In this context, representing exceptions in XML is particularly useful as it makes it easier for services to detect and add to incoming exceptions.

This approach of relaying exceptions to the client can be extended to handle metadata propagation. One example of metadata is information about billing from individual services. Metadata can be appended to normal data as it is passed between services. However, for services to process and exchange documents containing both the data and its metadata, a standard is required.

Finally, consider the issue of an unexpected delay occurrence in one of the services in a chain. The serial nature of the chain implies that the delay propagates through the chain, and all the way to the client. In the scenario where the client directly accesses each service, the client can abort the operation if a specified time-out period for a service expires. In that case, the client can opt for a substitute for the timed-out service. However, with nested calls, a global time-out is needed. This time-out is controlled by the client and may be communicated to every service in the chain. At any point, if a service takes longer than this global time-out, it aborts and returns an appropriate exception to the preceding service in the chain. The efficiency of chains can be further improved through data compression. Compressing data decreases transmission time, although it may be at the expense of increasing the processing time at the services.

The above analysis illustrates how client-coordinated service chaining in a decentralized setting requires deep involvement of the client. Although the coordination responsibilities of the client may be alleviated by simple interoperable service interfaces, it is not unlikely that they will hinder a wide-base adoption of the GIS Web Services model. The next section introduces static chaining by aggregate services as a service chaining alternative which hides all service chaining complexities from the client.

3.2.2 Static Chaining using Aggregate Services

With aggregate services (Figure 7) chaining is totally opaque to the client. The services appear as a single aggregate service which handles all the coordination of the individual services in a chain. Aggregate services bundle static (pre-defined) chains of services and present them to the client as one.

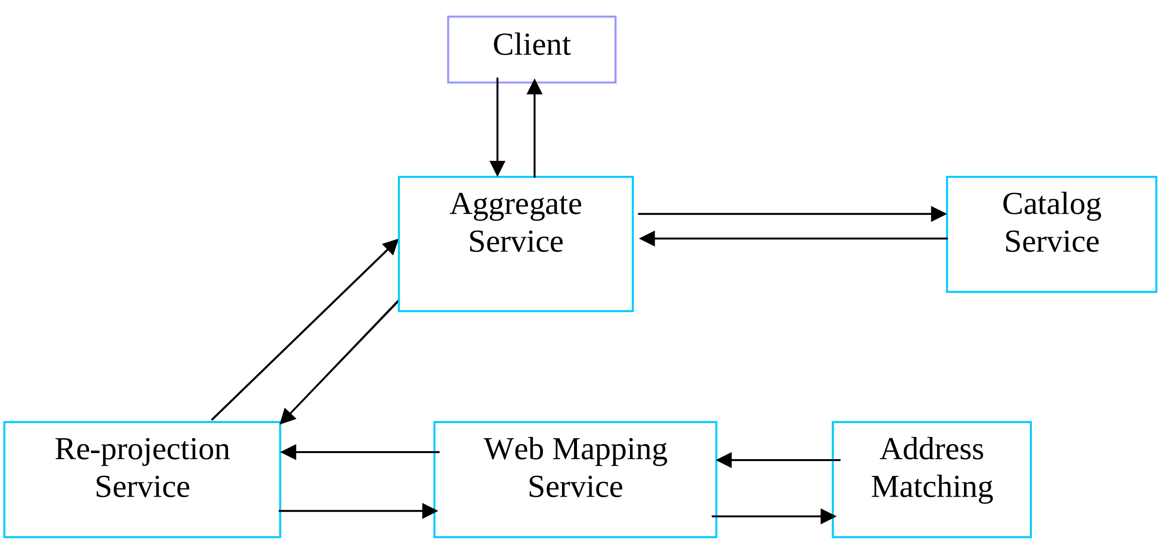

Figure 7. Static Chaining with Aggregate Services.

Despite their benefits, aggregate services have some drawbacks. By having a single access point to the chain, the client loses some of the flexibility and control over parameters of the individual services. In the example, the client has no control over the re-projection step. In fact, the client may not even be aware that the image is being re-projected. The invisibility of this step and the assumptions made by the aggregate service can be misleading to some clients. For this reason, clients should be able to differentiate between basic and aggregate services, which is why within OGC, service descriptions (or capabilities files) include a flag that indicates whether a service is an aggregate (cascading) service.

Fortunately, there is a more flexible and scalable alternative to aggregate services in a distributed environment where static binding of services and calls is often not efficient. The next alternative introduces workflow-managed service chaining with mediating services, which aims at achieving a balance between the totally transparent client-coordinated chaining and the totally opaque aggregate services chaining.

3.2.3 Workflow-Managed Service Chaining with Mediating Services

As seen in Figure 8, mediating services act as gateways to other services by coordinating between multiple services without necessarily storing any data of their own. Mediating services combine the simplicity of aggregate services with the flexibility and control inherent in client-coordinated service chaining. The client may execute the chain himself or just select a chain to be executed by the mediating service.

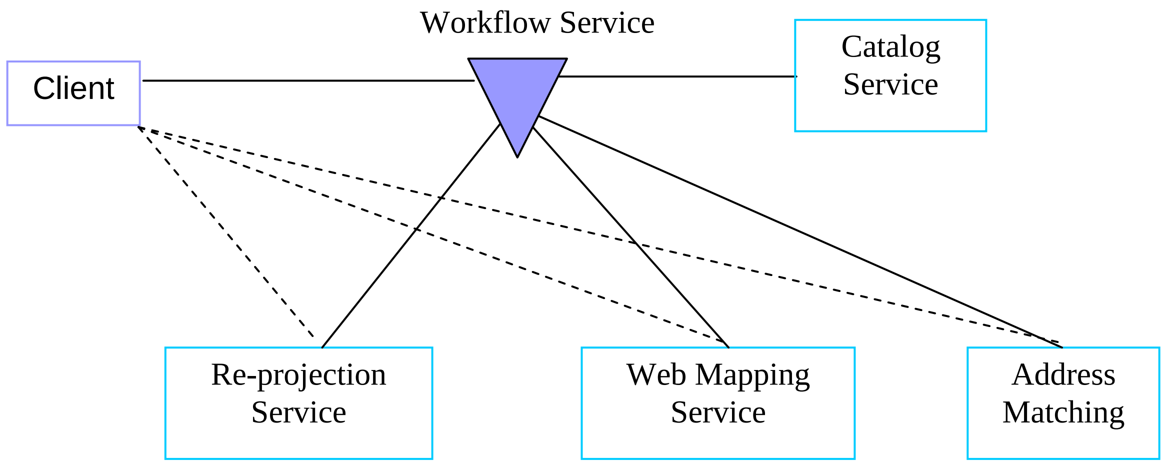

Figure 8. Workflow-Managed Chaining with Mediating Services.

The concept of mediating services is borrowed from the database arena (Litwin 1990). In a distributed database setting, mediating elements dynamically convert multi-database queries into smaller sub-queries that can then be dispatched to the various databases. The results of the sub-queries are then integrated by the mediating element and returned to the client. In the database literature, these mediating elements are also referred to as facilitators, brokers, and dispatchers (Wiederhold 1999).

Correspondingly, in a distributed geo-processing infrastructure, mediating services dynamically construct and manage chains of GIS Web services. Based on their client's requirements, mediating services determine appropriate data sources and services, retrieve and process the data, and then assemble the final response. In the process, a mediating service may consult with catalogs, search engines or meta-search tools known to it. It can also keep its own indexed lists of useful services, which are more likely to be biased towards certain providers and/or domains. For efficiency purposes, mediating services may also provide commonly used basic functions, such as format, coordinate or vector-to-raster conversions.

Mediating services use pre-specified client preferences to search for appropriate data and processing services that best meet their clients' requirements. Such preferences might include information about service time-outs, price ceilings, accuracy requirements, and maximum number of services chained. In some cases, the client may also wish to specify a preference for a particular service provider. The client may also impose a constraint that all services used in a particular session be supplied by the same provider, presumably to achieve certain efficiencies as well as monetary savings.

Mediating services can be considered as specialized versions of existing workflow and process management tools. The need for specialization is a consequence of the distinctive characteristics of GIS (Alameh 2001), especially the semantics associated with GIS data and the complexity of spatial queries. With the wide range of possible GIS applications and the different semantics needed in different fields, it is more likely that the internal mediating service rules will be tuned to specific application domains.

Therefore, in terms of market dynamics, we foresee the emergence of a variety of mediating service providers, which range in their "smarts" as well as the nature of their specialization. The need for domain-tuned services will constitute excellent market entry opportunities for third party players with significant expertise in a certain domain, but with no capabilities to single-handedly offer and maintain all the data and transformations needed for that domain.

As the number of services and catalogs available in an environment grows, there will be an increasing need for more sophisticated search-engine-like tools that can consolidate, organize and present information retrieved from various sources. Such tools may also provide interfaces through which users can pick services they need. Such tools can dispatch the users' requests to a variety of available catalogs, and then allow users to sort the results according to different criteria, e.g., price, quality or provider. As such, these tools are similar to popular online price comparison sites (e.g., metaprices.com or simon.com) which allow users to pick a category of items to compare (e.g., cds, books, electronics) and then return a list of items along with their prices, availability, special offers and reviews from various online shopping websites.

In summary, mediating services promise to minimize the complexity of service chaining while providing clients with solutions that are specifically tuned to their preferences. Mediating services also provide a client with a single point of contact for accounting and authentication, as well as error and metadata reporting.

4 Conclusion

It is becoming increasingly evident that there is a growing need for a GIS Web Services architecture (Abel D, Ooi B, Tan K, and Tan S 1998). This architecture will be especially beneficial for scientific research and engineering modeling as well as state and federal government settings, where tightly coupled hierarchical systems are unlikely to provide the desired breadth and flexibility. The web services model allows users in these settings to freely combine services to create customized solutions with minimal programming, integration and maintenance efforts. Such a model will also be a key enabler of GIS to extend beyond its traditional boundaries of mapping to embrace a broader community of users and wider scope of services. Critical to the success and sustainability of the distributed Web Services architecture is the issue of service chaining. After examining three service chaining options (client-coordinated, static and workflow-managed), we conclude that workflow-managed service chaining with mediating services combines the simplicity of static chaining with the flexibility and control of client-coordinated chaining.

Enabled by the advancements in web services in general, and by the on-going work of subgroups within OGC, the GIS Web Services architecture is rapidly manifesting itself. Various groups within OGC are working on service categorization (data, processing and registry services), encodings (SLD, GML), and service chaining (WMS/SLD/CPS, which is setting a precedent for service chaining in the web services environment). Within this work, general web services technologies have been critical: examples include WSDL for service description, UDDI for service discovery, SOAP for passing XML-encoded data, and IBM WSFL and MS XLANG for web service composition and process languages for orchestrating web services (OWS1 2002).

The next section provides some insights on how the GIS marketplace may be changing in the near future as a result of such advancements.

4.1 *A Changing GIS MarketPlace*

The unbundling of GIS systems into independently-provided interoperable components, and the delivery of subsets of GIS data to users on demand will lead to significant changes in the GIS marketplace. Figure 9 outlines a potential value chain for the future GIS marketplace.

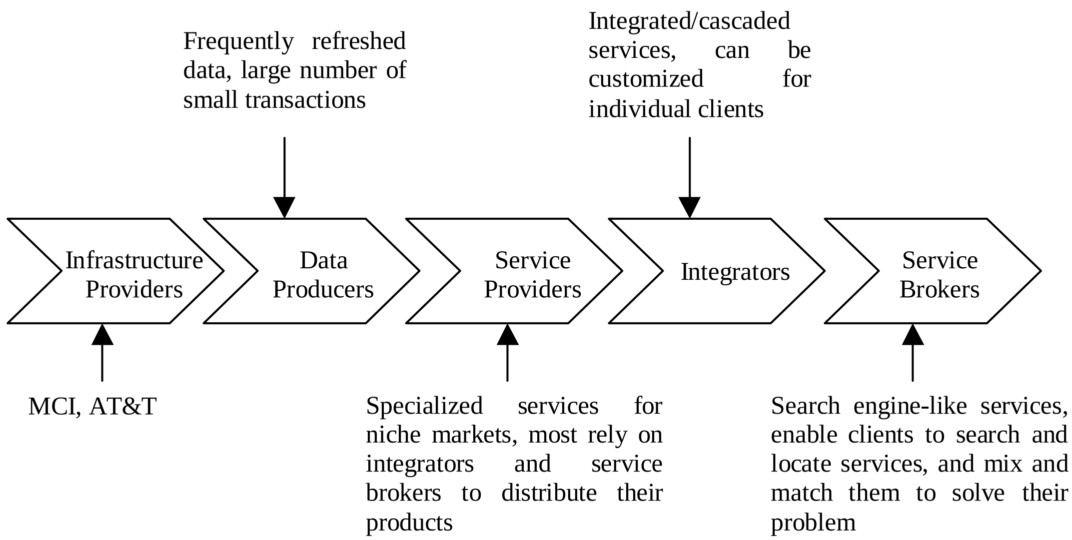

Figure 9. Potential Value Chain for the Future GIS Marketplace.

In this new distributed environment, the private sector as well as the public sector at the local, state and federal levels will all likely contribute in establishing and maintaining a national GIS infrastructure. The different players in the value chain will share different responsibilities according to their expertise. For instance, governmental agencies are well positioned to offer and maintain public data covering their areas of jurisdiction. National agencies such as NASA and NIMA can provide shared access to their data via interoperable interfaces. In the private arena, satellite imagery providers are likely to follow an e-commerce model for providing users with on-demand access to their huge repositories of data.

With the unbundling of GIS, it will not be necessary for players to build comprehensive systems in order to gain a share of the market. The new environment will open the door to small niche players to enter this market with application specific offerings that leverage their understanding of particular industries or processes. In this paper, we saw that the need for mediating services to coordinate service chaining will provide huge market entry opportunities for these new players. New opportunities may be available for some service providers to target niche markets in the cases when the back-end services are expensive, when service chaining requires specific domain expertise, or when the data provided is sensitive to local context and subcultures. Nonetheless, these opportunities will be limited by the availability of data/service repositories and catalogs in the market. Such players are likely to wait for enough services to become available on the market, and select partners from the players that provide them.

Finally, in terms of the reaction of traditional GIS systems providers in the face of the new competition, we expect them to adapt their business models by offering access to components of their systems through portal-like applications. Until now, the traditional players have been intentionally slow at aggressively developing applications for thin client applications in order to protect their systems. In order to compete, these players will leverage their established brand names as well as their connections with their current customers. However, in order to maintain their current investments in their clients, we anticipate that the traditional players will tune their services to better perform when coupled with their own clients.

Of course, the path to realizing the potential changes described in this section requires standards for data and metadata exchange in addition to well-defined simple service interfaces, both current key research topics within the GIS community.

5 References

Abel D, Ooi B, Tan K, and Tan S 1998 Towards integrated geographical information processing. *International Journal of Geographical Information Science* **12**(4): 353-371.

Alameh N 2001 *Scalable and Extensible Infrastructures for Distributing Interoperable Geographic Information Services on the Internet.* PhD Dissertation Submitted to MIT Libraries.

Bernstein, Philip 1996 Middleware: A model for distributed system services. *Communications of the ACM* **39**(2): 86-98.

Brodie, Michael 1993 The promise of distributed computing and the challenges of legacy information systems. In Hsiao D K, Neuhold E J, and Sacks-Davis R (eds) *Interoperable Database Systems (DS-5)* Amsterdam: 1-32.

Buehler K, and McKee L 1998 *The OpenGIS Guide: Introduction to Interoperable Geoprocessing and the OpenGIS Specification*. Waltham, MA, OpenGIS Consortium.

Doyle A 2000. *Web Map Server Interface Specification: OpenGIS Document 99-077r6.* Waltham, MA, OpenGIS Consortium.

Finch I, and Small E 1999 Information brokers for a web-based geographic information system. In Goodchild M, Egenhofer M, Fegeas R, and Kottman C (eds) *Interoperating Geographic Information Systems.* Boston, MA, Kluwer Academic Publishers:195-202.

Gunther O, and Muller R 1999 From GISystems to GIServices: Spatial computing on the internet marketplace. In Goodchild M, Egenhofer M, Fegeas R, and Kottman C (eds) *Interoperating Geographic Information Systems.* Boston, MA, Kluwer Academic Publishers: 427-442.

ISO 2001 ISO/TC 211 Geographic Information/Geomatics. *http://www.statkart.no/isotc211/*

Litwin W, Mark L, and Roussopoulos N 1990 Interoperability of multiple autonomous databases. *ACM Computing Surveys* **22** (3):267-293.

Nebert D 1999 Interoperable spatial data catalogs. *Photogrammetric Engineering and Remote Sensing* **65**(5): 573-575.

OGC 2002 *OGC Web Services Initiative 1 (OWS1) Baseline Documents Page* WWW document, http://ip.opengis.org/ows1/docIndex.html.

Wiederhold G 1999 Mediation to deal with heterogeneous data sources. In *Proceedings of Second International Conference*. Zurich, Switzerland: 1-16.