On Quality of Service and Geo-service Compositions

Richard Onchaga
International Institute for Geo-Information Science and Earth Observation (ITC)
P.O. Box 6, 7500 AA Enschede, The Netherlands
Tel: +31 53 487 4228, Fax: +3153 487 4575
onchaga@itc.nl

SUMMARY

Geographic information services (geo-services) are gaining popularity as an efficient and cost-effective framework for integrating geo-information systems, within and between enterprises, to enhance enterprise business processes. Geo-services can be chained dynamically to create customised geo-information and value-added services and therefore offer a business opportunity for disparate partners to collaborate and address the needs of consumers in the marketplace. Nonetheless, commercial exploitation of geo-service technology demands a computing platform that facilitates chaining of geo-services taking into account user requirements, both functional and quality of service (QoS), and that is capable of handling changes in the marketplace. We call such a platform a geo-service infrastructure. A geo-service infrastructure enables geo-services possibly offered by different enterprises and with potentially different capabilities to be discovered, composed and executed to obtain customized information and value-added services suited to user needs. In this paper we prescribe a geo-service infrastructure and define how it facilitates QoS-aware chaining of geo-services. We identify the basic components of the infrastructure and define their main functions. We also elaborate on the design principles upon which the infrastructure is based and explore how the infrastructure can be realized.

KEYWORDS: *Geographic information services, geo-services, service chaining, quality of service, geo-service infrastructure*

INTRODUCTION

In recent years there has been much activity in the geographic information industry centred on geographic information services (geo-services). Geo-services are loosely-coupled, open and interoperable geo-processing modules that can be invoked across a network by any compliant system to access process, and produce geographic information (ISO 19119, 2002). Accordingly, geo-services present a computing framework that is interoperable, flexible and extensible. Within the enterprise, geo-services are emerging as an efficient and cost-effective means of integrating enterprise geo-spatial resources and making them sharable enterprise-wide. Similarly, geo-services provide the vital glue between geo-information systems and other enterprise information systems that is necessary to enhance the efficiency and effectiveness of enterprise business processes. Between enterprises, the services offer a scalable and extensible framework for integrating disparate and heterogeneous geo-information systems across technology and organizational boundaries (Tait, 2005).

Besides the ability to foster intra- and inter-enterprise system integration, geo-services have the inherent capability to be chained together or with other services to realize more elaborate functionality and deliver value-added services. In service chaining, geo-services, possibly offered by different enterprises and with potentially different capabilities, are discovered, composed and executed in processes to obtain customized information and value-added services that are suited to user needs. Service chaining can be realized statically at design time or dynamically during run-time. In this paper we focus on dynamic chaining. Dynamic chaining of geo-services creates a business opportunity for disparate providers to collaborate and share knowledge, geo-spatial and computing resources and create customised geo-information and value-added services for consumers in the marketplace.

Conventionally, geo-services are chained considering only their functional capabilities. This approach is inadequate particularly in competitive marketplaces where: users seek customized information and value-added services suited for their specific needs; there are potentially many providers offering geo-services that are similar in terms of functionality but significantly different in quality of service (QoS) and cost characteristics; availability of services changes over time i.e. services can be deployed or deprecated anytime.

Clearly, there is business value in a chaining approach that considers user requirements on the one hand, and functional and QoS characteristics of available resources (geo-services) on the other hand, as operands to select geo-services and other resources that instantiate an optimal service chain for a given set user requirements. We call such an approach QoS-aware service chaining. From a market perspective, QoS-aware service chaining facilitates service differentiation and makes feasible pervasive and on-demand access to customised geo-information and value-added services (Alameh, 2003). From a systems' perspective, QoS-aware service chaining offers a framework for evolving flexible geo-information systems and location-based services that can adapt to changes in technology and user requirements (Di, 2004).

In this paper we prescribe a *geo-service infrastructure* – a distributed computing platform for QoS-aware chaining of geo-services. First, we motivate a definition of QoS in the context of service-oriented geo-processing. Second, we identify the basic components of the geo-service infrastructure and define their main functions. Also, the design principles upon which the infrastructure is based are elaborated, and third, we explore how the geo-service infrastructure can be realized building on advances in web service and other technologies.

FROM DATA QUALITY TO QUALITY OF SERVICE

The notion of quality of service (QoS) is broad and is used in many domains, often with different interpretations. Nonetheless, in any domain, a precise interpretation of QoS demands that one has a crisp understanding of the service(s) that is/are the subject of quality, where a service is defined as the externally observable behavior of a system in this case the service-providing entity. For example in networking, the basic service a network offers to its users is to move data packets between cooperating processes running on different end-nodes. Accordingly, network QoS mainly concerns the delay packets experience in the network, the probability that packets are lost during transfer, and the need that packets arrive at the receiving processes in the correct order and sequence as they were dispatched by the sending process.

In service-oriented geo-processing, a service-chain provides access to geo-information and or processing functions. Depending on the application context therefore, the emphasis of quality concerns can be on the information delivered, or on both the information and the operational characteristics of the ad-hoc application that is realized by that instance of the service-chain. Accordingly, QoS in service-oriented geo-processing concerns both the quality characteristics of deliverable geo- information, and the qualities associated with the collective behavior of one or more geo-services which make up the ad-hoc application that creates the information; the so-called ad-hoc application.

Figure 1 depicts a QoS meta-model model for services provided by chaining disparate geo-services. The figure shows that QoS comprises quality characteristics of deliverable geo-information and quality characteristics of the (ad-hoc) application that creates the information – both of which determine the utility of a service in given context. Quality of geo-information as used here subsumes notions of spatial data quality and encompasses quality aspects of the computational models realized by the geo-services that collaborate to process, transform and fuse the input data, and the uncertainty associated with the input data as well as the computational model(s) (Heuverlink, 1998; Drummond, 1991). Figure 1 also shows that the quality of an instance of a service chain comprises the collective effect of the individual quality capabilities of individual geo-services in the instance of a service-chain that realizes the application and the QoS delivered by the underlying communication network.

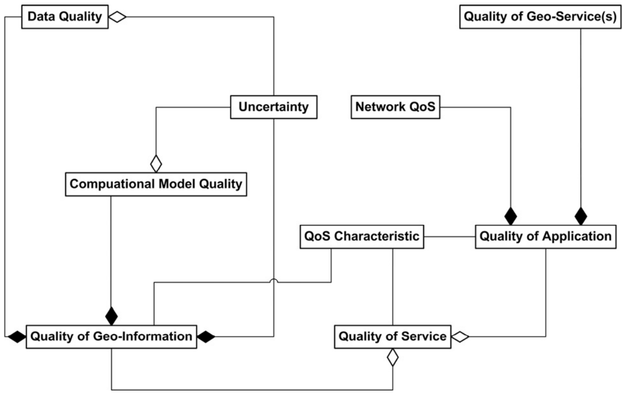

Figure 1 QoS meta-model

To illustrate the aspects of QoS introduced above let us consider the case of an orthophoto service. An orthophoto service provides online access to orthophotos. The process of orthophoto generation involves a number of operations that may include *imagery access*, *rectification*, *reprojection*, etc. Ideally, it is possible to provide orthophoto services by a coarse-grained monolithic service that implements all required processing operations or by a dynamically constituted chain of fine grained geo-services. For the purposes of this paper, we shall only elaborate on the latter further because if affords the flexibility and extensibility desired in service oriented architectures. Figure 2 shows some geo-services that are used by an orthophoto service and the functions they perfom. The figure shows that an orthophoto service can use instances of *mapping*, *imagery*, *rectification* and *reprojection* service.

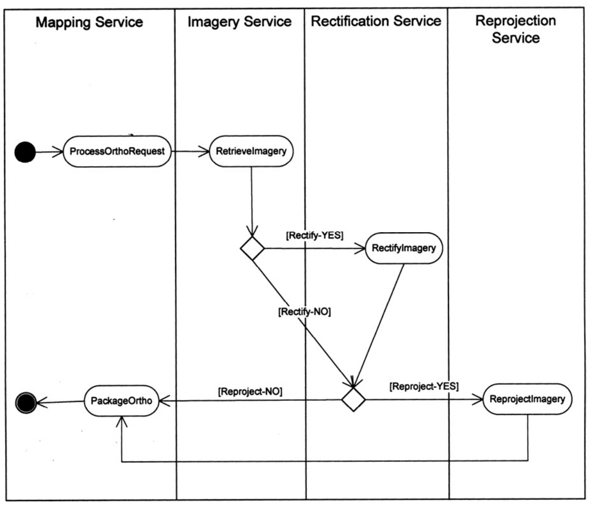

Figure 2 *Orthophoto service-chain*

Orthophoto products are widely used for example, as primary sources of data to populate GIS databases, to generate digital elevation models (DEMs), topographic maps, and thematic maps for a variety of applications e.g. engineering applications, planning, cadastre and titling, disaster response, etc. However, each application context has specific requirements that define the desired quality of an orthorphoto product. For example, for populating GIS data bases, the geometric integrity of the orthophoto product is a primary concern, which translates into demand for imagery of greater quality, and more rigorous computational models i.e. for rectification and reprojection services. In this case, for instance, the computational model implemented by the rectification service needs to take care of both relief and tilt displacements. In contrast, orthophotos for titling in rural areas do not demand high degree of integrity and therefore lower precision data and more approximate computational models can be employed.

The above concerns on information quality also persist in a service chaining, but in addition QoS characteristics of the service chain come into play. Once again, consider using an online orthophoto service to deliver orthophoto products for planning and also for disaster response. In the latter case, the delay and availability of the service are critical to the success of the time-critical search and rescue operations whereas in the former delay and availability requirements are less critical. Further, it is apparent that the delay a consumer will experience in either case will include processing delays for executing each of the requisite operations and the cumulative transmission delay experienced by messages in the communication network.

THE GEO-SERVICE INFRASTRUCTURE

The geo-service infrastructure is the distributed computing platform for QoS-aware chaining of geo-services. The infrastructure facilitates discovery and selection of geo-services, composition of selected geo-services and management of the execution of resultant service-chain instance all on the basis of user-specified QoS requirements and quality characteristics of individual geo-services. The activities involved in QoS-aware discovery, selection and composition of geo-services and subsequent chain execution management are all transparent to the user (or client application). As such the geo-service infrastructure shields the users from the complexities of QoS-aware service chaining. Similarly, the infrastructure relieves geo-services developers and deployers of the complex tasks of QoS-aware geo-service selection, composition and service-chain execution management and thus simplifies geo-service development and deployment.

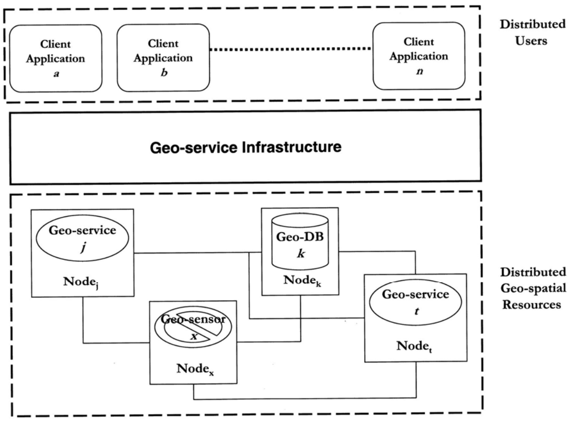

Figure 3 Geo-service infrastructure

The geo-service infrastructure provides QoS transparencies to users and providers of geo-services. Figure 3 shows that the geo-service infrastructure sits between the client applications (users of geo-services) and the geo-spatial resources (geo-services, geo-databases, geo-sensors, etc.) in distributed nodes that are interconnected by communication networks. The geo-spatial resources in distributed nodes and the communication networks together constitute a distributed resource platform (DRP) for service-oriented geo-processing. The client applications, geo-service infrastructure and distributed geo-spatial resources realize three architectural layers which we call *user*, *application* and *resource levels* respectively.

Ideally, beside QoS-enabling services, the geo-service infrastructure will encompass distributed messaging service, service-chain definition and enactment services, and conventional registry services but we restrict our discussion to the QoS-enabling services in this paper. QoS-enabling services can be deployed in a centralized or distributed architecture. In either case, there are advantages and disadvantages. For example a centralized architecture is easier to deploy but is less scalable. In contrast, a distributed architecture is more scalable but demands more complex coordination among collaborating QoS-enabling services. These architectural concerns will not be elaborated further in this paper. We call the QoS-enabling services QoS mechanisms three categories of which are distinguished: QoS specification mechanisms; QoS control mechanisms; and, QoS management mechanisms. In the following sections we present these basic categories of QoS mechanisms and outline the fundamental principles that underlie development, deployment and successful operation of the QoS mechanisms.

QoS specification mechanisms

QoS specification concerns requirements and QoS (management) policies associated with the requirements. End users of a service have certain requirements or expectations on a service. These are called user requirements. For instance in the orthophoto service example previously described, in a disaster response scenario, it may be required that orthophoto products are delivered within some specified delay limit and that services should be available throughout the entire period of search and rescue program. Typically, user requirements are expressed in terms that are meaningful to the user and for the example, the user requirements may be specified as constraints on *interactivity* and *dependability* e.g *interactivity* $\leq$ *5 sec* and *dependability* $\geq$ *99%*. Given the user requirements, corresponding application level QoS requirements can be derived and suitable service providers identified. The derived QoS requirements are also the basis on which appropriate geo-services at the resource level are selected to instantiate a service-chain.

QoS policies are guidelines that specify the way a service provider shall control, manage and report achieved service levels including cost of service and compensation to the users, where appropriate, in the event of the provider failing to maintain agreed service levels over the period of service delivery. QoS requirements and QoS policies together constitute a service level agreement (SLA).

The preceding discussions point to the existence of varying notions of QoS at the different levels of the architecture (Figure 3). First, at the user level, user requirements capture the QoS needs or expectations of users. At the level of service (or application) QoS is specified as QoS requirements and typically concerns aggregations of atomic resources. Lastly, at the resource level, QoS concerns QoS capabilities of individual resources. Typically, the notions of QoS at the different levels are specified in terminology appropriate in each level and requirements across the different levels in general exhibit many-to-many type of relationships. Accordingly, the following QoS specification mechanisms are required:

- QoS mapping mechanisms that automatically map requirements across the different architectural levels. For example user level *interactivity requirement* can be mapped onto application level *reliability*, *availability*, *delay*, etc. Further, application level delay can be mapped onto resource level *performance*, *network bandwidth*, etc., of individual geo-services and network links.
- Discovery mechanisms that query QoS capabilities and state of resources at the DRP.

- QoS composition mechanisms determine the set of geo-services and (other resources) from available resources that optimize a given service-chain with respect to specified requirements.

QoS control and QoS management mechanisms

In essence QoS specification mechanisms achieve QoS prediction and establishment (ISO/IEC 13256, 1998) that are activities which happen before actual execution of a service chain starts. During the period of actual execution of a service chain i.e. the operational phase of QoS, QoS control and QoS management mechanisms are required. In principle, QoS control and QoS management mechanisms are different categories of mechanisms but here we discuss them together because both types of mechanisms are used in the operational phase of QoS, the main difference being the time span over which they they operate.

Typically, QoS control mechanisms operate over the time span of an atomic interaction. During the execution of a service chain, many interactions will take place, sequentially or in synchronously, between entities towards delivering a desired service. Each atomic interaction will involve a pair of entities in the roles of service user and provider respectively. The purpose of QoS control mechanisms is to ensure each atomic interaction proceeds successively and in case of a failure to provide an appropriate QoS alert to the relevant QoS management function. As an illustration, consider the case of the orthophoto service again. A mapping service may request a rectification service to reproject orthophoto imagery within a given deadline. Now supposing the reprojection services becomes unavailable for some reason, QoS control mechanisms should raise an alert therefore allowing an alternative reprojection service to be employed.

QoS management mechanism essentially perform the same functions as QoS control mechanisms but operate at a longer time span i.e. the duration of execution of an entire service-chain. QoS management mechanisms monitor the execution of a geo-service chain and enable QoS maintenance. The mechanisms are also responsible for informing end-users on QoS violations. QoS maintenance concerns monitoring the execution of a geo-service chain and incase of impending violation of QoS, re-assign succeeding operations to alternative geo-services to maintain QoS targets. For example consider that the operations *retrieveimagery*, *rectifyimagery* and *reprojectimagery* (Figure 2) are to be performed sequentially in an instance of a service chain to deliver an orthophoto service. Now assume that operation *retrieveimagery* is successfully executed and *rectifyimagery* is invoked. If for some reason, rectification service becomes unavailable, therefore terminating the *rectifyimagery* operation, the QoS control mechanisms raise a QoS alert as before. Following the alert, QoS management mechanisms will have to reassign the operation *rectifyimagery* to a different rectification service and reschedule the *reprojectimagery* operation so as to achieve contracted service levels otherwise an appropriate alert is sent to the end-user of impending QoS violations.

It is apparent that QoS management mechanisms depend on QoS control mechanisms (also QoS composition and query mechanisms) to achieve their functions. Further, both QoS control and management mechanisms in the geo-service infrastructure must necessarily collaborate with corresponding mechanisms at the resource level to realize their functions.

QoS principles

To develop, deploy and successfully operate a geo-service infrastructure that provides adequate QoS transparencies to users and providers of geo-services, a number of basic design principles are employed. These principles, which we call QoS principles, are in addition to the principles of the service oriented architecture (SOA) that geo-services are assumed to follow. We define three QoS principles: transparency principle; integration principle; and, separation principle.

- Transparency principle – the transparency principle states that the activities of QoS-aware service chaining should be transparent to service consumers. The essence of the transparency principle is that the complex procedures associated with QoS-aware service chaining should be hidden from the service consumers as much as possible. By service consumer we mean any entity that requires some service provided by another entity in the framework of service-oriented geo-processing. Service consumers will therefore include end-users and service providers.

- Integration principle – the integration principle states that QoS should be configurable, predictable and maintainable across all architectural levels. In essence, the integration principle considers QoS as being end-to-end i.e. the QoS perceived by an end-user is the collective result of QoS effects across all architectural layers and includes the QoS delivered by every entity in the process of delivering a service. It is on the basis of this principle that QoS control and management mechanism in the infrastructure and those on the lower level resource layer collaborate to control and manage QoS across the architectural levels.
- Separation principle – the separation principle states that the functions that enable interactions between geo-services (and other infrastructure services), the functions for QoS control, and the functions for QoS management realize distinct architectural activities and should therefore be separated in architectural frameworks. The essence of the separation principle is to achieve a clean separation of concerns in geo-service architectures and enable development of an evolvable and scalable infrastructure.

IMPLEMENTATION CONCERNS AND RELATED WORK

A geo-service infrastructure is an abstract notion that should be realized by one or more concrete entities that collaborate to deliver the functionality ascribed to the infrastructure. A basic building block for the geo-service infrastructure is the ability of resources that are shared and consumed i.e. services and artifacts at the resource level, to be self-describing both in terms of their functional and QoS capabilities. The notion of a geo-service infrastructure builds on this capability. It is apparent therefore that standardized resource descriptions that include QoS information are a prerequisite to realizing a meaningful QoS-aware chaining of geo-services.

The Grid and web services are alternative technologies for implementing geo-services but the thrust of geographic information industry efforts has centered on web service technologies. In a web service environment, the web service description language (WSDL) is the de facto standard for describing web services. Earlier versions of WSDL i.e. WSDL 1.1 provided no support for specifying QoS information of services. This problem will hopefully be fixed in the proposed WSDL 2.0 which proposes a Features and Properties component architecture that is potentially extensible to incorporate QoS information. Furthermore, parallel efforts outside of World Wide Web Consortium (W3C) for example efforts on Web Service Offerings Language (WSOL) (Tosic, et al., 2003), Web Service Level Agreement (WSLA) from IBM Corporation (Ludwig et al., 2003), and other similar efforts are addressing the issue providing QoS information and it is anticipated that a convergence in these efforts can be reached to provide a standard way of capturing QoS information in service descriptions.

	Informational	Operational
User Level	— Accuracy (reliability) — Fidelity	— Interactivity — Dependability — Price
Application Level	— Information quality e.g. freshness, geometric integrity, accuracy, etc.	— Delay — Reliability — Security — Cost

Table 1 Quality model (adapted from Onchaga, 2005)

Complementary to the ability to describe QoS of services is the need that users are able to specify their requirements. This can be achieved relatively easily by extending client interfaces or by employing agents that detect and maintain the QoS context of a client. In general, user requirements are specified in terms that are meaningful in their application context. To promote interoperability within a domain there is need therefore for a quality model defining commonly used quality characteristics, the corresponding value domains and the constraints that can be expressed over them. However for interoperability across domains and across architectural layers a QoS mapping scheme becomes necessary. In our earlier work (Onchaga, 2005) we have elaborated on these issues – specifically we defined an extensible QoS model for services delivered through dynamic composition of geo-services and demonstrated a QoS mapping scheme based on the technique of translation tables. Table 1 illustrates a generic and extensible "two-dimensional" multi-level quality model; the two dimensions correspond to informational and operational QoS requirements that are typical in service-oriented geo-processing. Levels correspond to the architectural level to which a characteristic applies e.g. interactivity is a user level operational requirement, etc. For brevity, Table 1 only shows two levels. The mapping scheme we have defined can be implemented by a software entity to realize a QoS mapping mechanism.

Lastly, mechanisms are required to select, compose geo-services and successfully execute resultant service chains to meet user requirements. This is the realm of QoS composition, control and management mechanisms. Moreover QoS control and management mechanisms at the infrastructure level also depend on the availability of resource QoS control and management mechanisms.

A significant amount of literature is available dealing with techniques and tools for managing quality of web services. Menasce and Almeida (2000, 2002) for example, detail approaches to capacity planning and define performance and workload models that can be used to deploy and configure web services to provide predictable QoS. Within the framework of achieving web services with predictable QoS, Shan et al. (2002) developed and tested load balancing techniques using priority of in-coming requests in a web cluster environment. Cardellini et al. (2002) define and test in a web cluster environment also in the framework of making predictable QoS of web services. Clearly, QoS mechanisms at the resource level are becoming available.

Many authors have also investigated QoS in the context of web service compositions. For example, Cardoso et al. (2002) have developed a predictive QoS model that makes it possible to compute the QoS for workflows automatically based on QoS attributes of atomic tasks. They also describe an algorithm to compute, analyze and monitor workflow QoS metrics. Similar work is reported by Zeng et al. (2003) who advocate a global planning approach for optimal selection of component services for chaining in a process and formulate the service selection activity as an optimization problem. Our work can also benefit from the work of Chandrasekaran et al. (2003) who describe a service composition and execution tool (SCET) for dynamic composition of service processes. They also describe and test a simulation methodology that can be used to test the performance of a web process. Similar work is reported by Casati et al. (2002) who have developed a system they call *eFlow* that support specification, enactment and management of composite e-services. The main goals of *eFlow* are to facilitate flexible and transparent adaptation of composite services in terms of QoS and also and dynamic modification of web processes. These studies indicate a growing number of approaches and tools that facilitate automatic composition and management of web processes. These tools to different degrees realize QoS composition, control and management mechanisms in the framework of a geo-service infrastructure and can be adapted and extended for use in the geo-service infrastructure.

A subject that requires further development however, is QoS composition. In the case of the geo-service infrastructure both constraints on the operational characteristics of available services and constraints on required information must be considered. As such, more fitting QoS composition algorithms that considers both informational and operational requirements are require to fully meet the requirements of users. Appropriate algorithms must take into account concerns like quality of input data, quality of computational model employed by a geo-service and the uncertainty associated with both model and data, in addition to and before operational characteristics.

CONCLUSIONS AND OUTLOOK

This paper prescribed a geo-service infrastructure that is a distributed computing platform for QoS aware chaining of geo-services. QoS-aware service chaining is needed for pervasive and commercial exploitation of geo-services and is especially desirable in competitive geographic information marketplaces where the number of available services and user requirements are dynamic. A definition of QoS in the context of service-oriented geo-processing was motivated and basic components of the infrastructure defined. The paper focuses on QoS-enabling components which we call QoS mechanisms. The core QoS mechanisms are introduced and defined. We also present a minimal set of QoS principles that are guidelines for successful development, deployment and operation of a geo-service infrastructure. Advances, especially in web service technologies, promise to provide the necessary capabilities for implementing a practical geo-service infrastructure. Nonetheless, more fitting QoS composition mechanisms are required that can resolve both informational and operational QoS requirements that are typical in service-oriented geo-processing. These algorithms are the focus of our on-going work.

BIBLIOGRAPHY

Alameh, Nadine S., 2003 Chaining of geographic information web services. In IEEE Internet Computing, Vol. 7 (5), pp 22 – 29.

Cardellini, V., Casalicchio, E., Colajanni, M., Mambelli, M., Enhancing a web-server cluster with QoS-mechanisms. In Proceedings of the IEEE International performance, computing, and communications conference (IPCCC 2002). Phoenix, Az. April 2002

Cardoso, J., Miller, J., Sheth, A., Arnold, J., (2202) Technical report: Modelling quality of service for workflow and web service processes. LSDIS Lab, Department of computer science , University of Georgia, USA.

Casati, F., Ilnicki, S., Jin, L., Krishnamoorthy, V., Shan, M., 2000 Technical report: Adaptive and dynamic service composition in *eFlow*. HP Laboratories, Palo Alto, USA.

Chandrasekaran, S., Miller, John A., Silver, Gregory S., Arpinar, B., Sheth, Amit P., 2003 Performance analysis and simulation of composite web services. In International Journal of Electronic Commerce & Business Media, Vol. 12 (2), pp 120 – 132.

Di, Liping, GeoBrain – A web services based geospatial knowledge building system. In Proceedings of NASA Earth science technology conference (CD-ROM). Palo Alto, CA. June 22 – 24, 2004.

Drummond, Jane E., 1991 Determining and processing quality parameters in geographic information systems. PhD Thesis, University of New Castle upon Tyne.

Heuvelink, Gerald B.M., 1998 Error propagation in environmental modelling with GIS. Taylor and Francis

ISO/IEC.13256, 1998 ISO/IEC 13236 Information Technology – Quality of service: Framework.

ISO/TC 211, 2002 ISO/DIS 19119 Geographic Information – services

Ludwig, H., Keller, A., Dan, A., King, Richard P., Franck, R., 2003 Web service level agreement (WSLA) language specification, Version 1.0. IBM Corporation.

Menasce D.A., Almeida V.A.F., 2000 Scaling for E-Business – Technologies, models, performance and capacity planning. Prentice Hall.

Menasce D.A., Almeida V.A.F., 2002 Capacity planning for web services: Metrics, models and methods. Prentice Hall PTR

Onchaga, Richard, Towards quality-aware composition of geo-services. In Proceedings of GSDI-8 Conference. Cairo, Egypt. April 16 – 21, 2005.

Shan Z., Lin C., Marinescu D.C., Yang Y., 2002. Modeling and performance analysis of QoS-aware load balancing of web-server clusters. Computer Networks,(40): 235-256.

Tait, Michael G., 2005 Implementing geoportals: applications of distributed GIS. In Computers, Environment and Urban Systems, Vol. 29, pp 33 – 47

Tosic, V., Patel, K., Pagurek, B., 2003 Research report SCE-03-21: Reusability constructs in the Web Service Offerings Language (WSOL). Department of systems and computer engineering, Carleton University, Ottawa, Canada.

Zeng L., Benatallah B., Dumas M., 2003 Quality driven web services composition. Proceedings of Twelfth World Wide Web Conference (WWW2003). Budapest, Hungary.