Towards Quality-aware Composition of Geo-services

Richard ONCHAGA, The Netherlands

Key words: Geographic information services (geo-services), service chaining, quality of service (QoS), QoS provisioning, geo-service infrastructure

SUMMARY

Geographic information services (geo-services) are gaining prominence as a framework for providing on-demand access to geographic information and value-added services. Loosely-coupled, modular and interoperable geo-services are discovered and chained on-demand to deliver geographic information and realize value-added services. In conventional approaches, chaining is accomplished considering only functional capabilities of geo-services. However, more valuable and effective service-chains can be realized by considering both functional and quality of service (QoS) capabilities. The latter type of chaining is called quality-aware service chaining. In quality-aware service chaining, disparate services are discovered and composed based on both their functional and QoS capabilities and subsequently executed in such a way as to provide services that comply with user requirements. Clearly, quality-aware composition of geo-services requires an effective QoS provisioning infrastructure in geo-service architectures to facilitate quality-aware chaining of geo-services. Frameworks have been defined that can be used to design, develop, and deploy effective QoS provisioning infrastructures for geo-service chaining. This paper expounds on QoS specification and QoS mapping, which are fundamental QoS provisioning functions, in the context of a QoS provisioning framework. The orthophoto service is used as a vehicle to derive QoS requirements and integrate the framework with service-oriented geo-processing. The technique of translation tables is used to realize QoS specification and mapping across architectural levels.

Towards Quality-aware Composition of Geo-services

Richard ONCHAGA, The Netherlands

1. INRODUCTION

Geographic information services (geo-services) have in recent years gained prominence as a new paradigm for developing and deploying distributed geo-information systems. A salient feature of geo-services is their ability to be discovered and chained dynamically, which positions them as a novel framework for evolving flexible geo-information systems and for providing on-demand access to geographic information and value-added services (Alameh, 2003).

Dynamic chaining leverages loosely-coupled, modular and interoperable geo-services to deliver geographic information and realize value-added services on-demand. In conventional approaches, chaining is accomplished considering only functional capabilities of geo-services. Nonetheless, more valuable and effective service-chains can be realized by considering both functional and quality of service (QoS) capabilities. We call this type of chaining quality-aware service chaining. In quality-aware service chaining, disparate services are discovered and composed based on both their functional and QoS capabilities to offer services that comply with user requirements.

QoS concerns the characteristics of a service that determine its utility in an application context. Accordingly, QoS in the context of geo-services and geo-service chaining comprises desirable qualities on geographic information delivered by a chain of geo-services and the qualities associated with the collective behavior of the geo-services (and other services) that create the service chain (Onchaga, 2005). Clearly, quality-aware composition of geo-services requires an effective QoS provisioning infrastructure in geo-service architectures to facilitate quality-aware chaining of geo-services. Frameworks are appearing that can be used to design, develop, and deploy effective QoS provisioning infrastructures for geo-service chaining. The frameworks are a base on which to design, develop, and deploy effective QoS provisioning infrastructures that apply user requirements and descriptions of available geo-services as operands to discover, compose and execute an appropriate chain of geo-services that delivers a desired service. This paper expounds on QoS specification and QoS mapping in the context of a QoS provisioning framework for service-oriented geo-processing we have previously proposed (Onchaga, 2005), in which quality-aware composition of geo-services is facilitated by the *geo-service infrastructure*. The orthophoto service is used as a vehicle to elaborate and demonstrate QoS specification and QoS mapping across different architectural levels.

The rest of the paper is organized as follows: In Section 2 we outline the QoS provisioning framework for service-oriented geo-processing; Section 3 introduces the orthophoto service; Sections 4 and 5 respectively present QoS specification and QoS mapping in the context of the orthophoto service; Section 6 concludes the paper.

2. QoS PROVISIONING FRAMEWORK

The QoS provisioning framework we use in this paper is outlined in (Onchaga, 2005). The framework defines the principles, concepts and mechanisms required to specify, develop and deploy a *geo-service infrastructure* for efficient quality-aware composition of geo-services. The *geo-service infrastructure* is the distributed computing environment that provides the functions necessary for quality-aware composition of geo-services. The QoS framework supports a layered architecture for QoS provisioning. Figure 1 shows the different layers (levels) of the architecture. The figure shows that the *geo-service infrastructure* sits between client applications, which end-users apply to access and exploit distributed (geo-) resources, and the distributed resources that users seek to exploit.

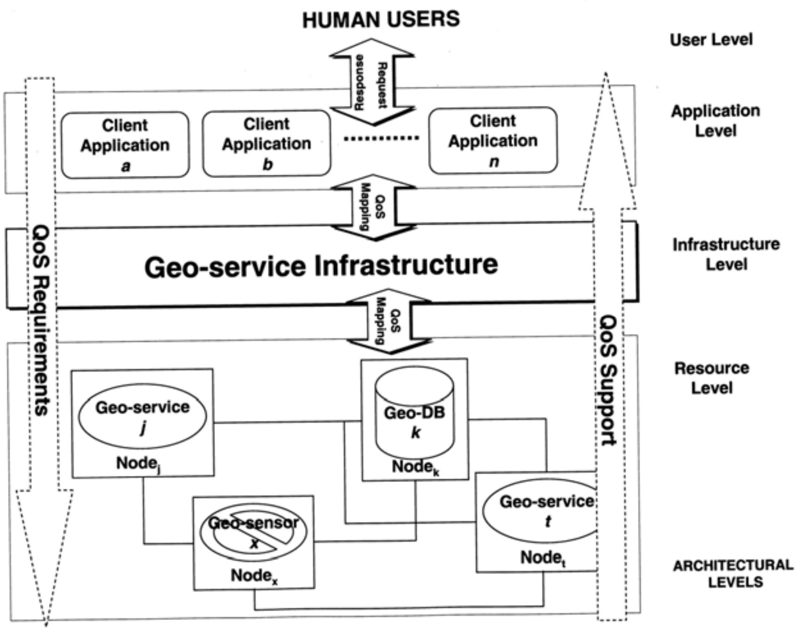

Figure 1: Architectural levels

QoS abstractions at the different levels are related but differ strongly in their interpretation (Siqueira & Cahill, 2000). The user level concerns QoS perceptions of end-users. The basic concerns at user level are service predictability and fitting cost-performance trade-offs that are essential for user-satisfaction (Widya et al., 2001). QoS abstractions at the application level concern proper operation of a service to achieve user-satisfaction. QoS abstractions at the application level, also called *user requirements*, are specified using concepts that are meaningful in the problem domain, and are applied to derive meaningful domain-independent QoS requirements called *application requirements*.

At the resource layer QoS concerns individual resources. Therefore QoS abstractions at the resource layer, also called *resource requirements,* are domain and technology dependent. Part of the functions of the geo-service infrastructure is to map between the different QoS abstractions. In the QoS provisioning framework, QoS specification and mapping are realized at the interfaces of interacting entities.

3. ORTHOPHOTO SERVICE

The example we use to expound on QoS specification and QoS mapping is the orthophoto service. The orthophoto service delivers orthophoto products that are customized to user requirements – on-line and on-demand. The term *orthophoto product* is used here in a general sense to refer to the result of any pre-processing or value-adding activity on orthophoto imagery. An orthophoto is the result of a recfitication process that corrects for tilt and or relief displacements inherent in aerial photography. The attracive features about orthophotos are that imaged features appear on an orthophoto in their natural form and are in their correct planimetric positions. Therefore imaged features are easy to identify on an orthophoto and measured distances, areas, volumes, etc., correspond well to their ground equivalents.

Orthophotos products are therefore widely used for example, as primary sources of data to populate geo-spatial databases, to generate digital elevetion models (DEMs), topographic maps, and thematic maps for a variety of applications e.g. engineering applications, real estate, planning, cadastre and titling, etc. Accordingly, orthophoto products have a high market potential which makes them ideal for dissemination in electronic marketplaces.

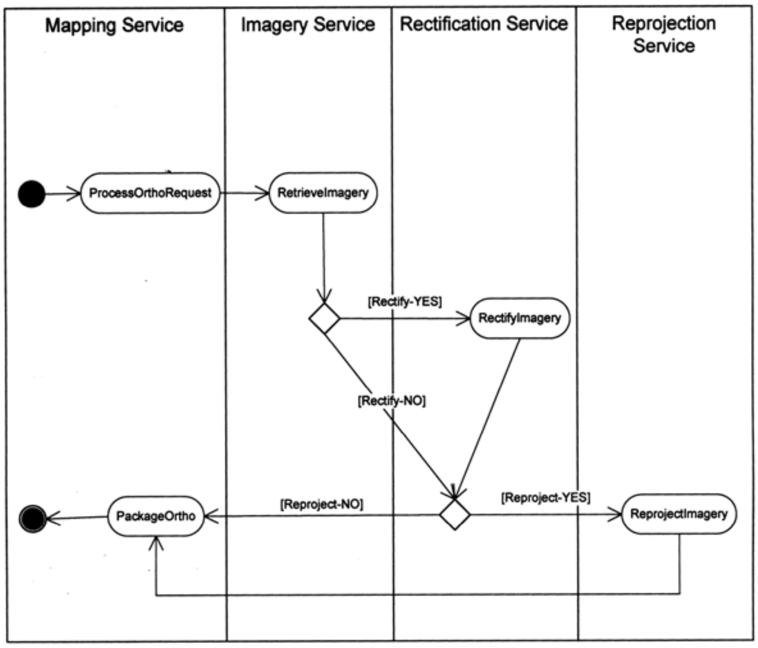

Figure 2: Geo-services and activity for orthophoto service

In a service oriented architecture (SOA), orthophoto services are realized by chaining disperate geo-services. Figure 2 is a simplified UML activity diagram showing geo-services that play a part in an orthophoto service and the activities they execute. The geo-services are briefly described below.

- The *mapping service* provides the front-end to the orthophoto service. The client interacts with the *mapping service* to make requests for orthophoto products and it is the responsibility of the *mapping service* to interpret incoming requests and determine the product specifications for incomming requests. The *mapping service* is also responsible for *packaging* i.e creating an appropriate rendering, formating and delivering the orthophoto products.

- Primarily, the *imagery service* provides access to imagery (digital aerial photos or orthophotos) in imagery achives.
- The *rectification service* transforms a tilted aerial photograph into its vertical equivalent. The *rectification service* may also implement algorithms to correct for relief distortions and generate a rectified orthophoto.
- Finally, the *reprojection service* is an image coordinate conversion service that transforms imagery from one coordinate system to another. Given an imagery in one coordinate system and a target coordinate system a *reprojection service* gives as output corresponding imagery in the target coordinate system.

4. QoS SPECIFICATION

The first step to quality-aware composition of geo-services is abstraction of user requirements. A user requirement is a quantifiable aspect of quality that is desired of a service or that is necessary for successful consumption of a service to occur. A set of user requirements that apply in a given problem domain constitute a QoS category for that problem domain (ISO/IEC, 1998). On the basis of user requirements, pertinent QoS requirements are derived and applied in QoS negotiation schemes to establish a service level agreement (SLA) that will enable the service provider to configure services that meet user requirements. An SLA will therefore include specifications on expected service levels and specifications on how service levels are monitored and reported including cost information.

The orthophoto service exemplifies service-oriented geo-information processing. In general therefore, orthophoto service delivery will face, to varying degrees, QoS challenges due to the relatively large volume datasets and compute-intensive processing tasks that characteritize geo-information processing. On the one hand, large volume datasets are bandwidth intensive. This is exercebated in *chatty* interactions involving extensive data exchange. Compute intensive tasks on the other hand are demanding on local computing resources. Table 1 presents user requirements and application level requirements relevant to the orthophoto service. The requirements in Table 1 are by no means exhaustive. Both the user level and application level requirements are further categorised into *informational* and *operational* requirements depending on whether they concern delivered geo-information or the operational characteristics of the service (Onchaga 2005).

The user requirements are:

- *Accuracy* is the fitness for use of an orthophoto product for intended use. Accuracy can be refined into more concrete sub-elements e.g. freshness, positional accuracy, etc.
- *Fidelity* concerns presentation or delivery characteristics of an orthophoto product with respect to user requirements and capabilities, for example the same orthophoto product when delivered to a networked workstation must be formatted differently from when it is delivered to a PDA, etc.

- *Interactivity* is the swiftness with which an orthophoto product is made available upon request.
- *Dependability* concerns the extent to which an orthophoto service is available and able to execute and deliver desired orthophoto products when needed and according to expectations. Dependability may also relate to the trustworthiness of an orthophoto service as perceived by the user based on experience or on advice of trusted third-party agents.
- *Price* is the cost charged for consuming a service against the perceived value of the service.

	Informational	*Operational*
User requirements	— Accuracy — Fidelity	— Interactivity — Dependability — Price
Application QoS requirements	— Information quality	— Delay — Reliability — Security — Cost

Table 1: Quality model for orthophoto service

Application QoS requirements include:

- *Information quality* is the external quality of information content (e.g an orthophoto) according to user requirements. Information quality is specified along one or more standard quality elements or sub-elements (ISO, 2002).
- *Delay* concerns timing aspects of an orthophoto service. Delay is the duration between the instant a request is submitted and the instant a response to the request is received.
- *Reliability* is the extent to which a service is available and, once invoked, executes correctly without interuption.
- *Security* concerns the level of security desired in a service and the corresponding mechanisms and protocols that realize secure services.
- *Cost* is the amount of money charged for consuming a service.

Orthophoto products are used in many domains. Typically, the different problem domains will require different levels of service. This is illustrated in Table 2 which categorizes potential application domains for the orthophoto service into four generic types ranging from *on-line professional* to *off-line non-professional*. By an on-line application we mean that the deliverables of the orthophoto service are consumed on-line i.e. on the fly. In contrast, off-line applications download the deliverables for off-line use .The *professionalism* of an application refers to the demands the application places on the precision and quality of deliverable products. Therefore an application categorized as being *professional* requires highly accurate orthophoto products and vice-versa.

From a providers' perspective, the four types of applications present four different QoS classes. Table 3 shows the four types of applications (QoS classes) with associated examples of problem domains. The classes are labeled; *platinum*, *gold*, *silver* and *bronze* corresponding to *on-line professional, off-line professional, on-line non-professional* and *off-line non-professional* types of application respectively where platinum.≥.gold.≥.silver.≥.bronze. Similarly, the service levels in Table 3 have the following semantics: high ≥ medium; super ≥ good; A1 ≥ A2, D1 ≥ D2 and the symbol "≥" reads as "better than". Service providers can deploy and publish orthophoto services that conform to one or more of the four classes. Cost is not shown in Table 3 as a user requirement because we believe that it is implicit and is part of the SLA for each class. However, different providers may cost similar services i.e. services that fall in the same class, differently.

Application type	Examples
On-line professional	Mission, command & control, disaster response, collaborative site selection, modeling applications, etc.
On-line non-professional	Location-based services e.g. fleet management, education and research, etc.
Offline professional	Primary data acquisition, engineering mapping, planning applications, etc
Offline non-professional	Educational applications, rural boundary delineation, etc

Table 2: Generic categorization of orthophoto applications

Application type/class	Accuracy	Fidelity	Interactivity	Dependability
On-line prof. (Platinum)	high	super	High	D1
Off-line prof. (Gold)	high	"don't care"	Medium	D1
On-line non-prof. (Silver)	medium	good	Medium	D2
Off-line non-prof (Bronze)	medium	"don't care"	Medium	D2

Table 3: Application types as QoS classes and associated service levels

5. QoS MAPPING

QoS mapping facilitates mapping between requirements at different QoS levels. Typically, requirements exhibit many-to-many relationships across levels. To map between requirements at different levels, the many-to-many relationships have to be resolved into less complex relationships. A mapping approach that utilizes translation tables has been proposed and applied in literature (Siqueira & Cahill, 2000; Widya et al. 2001) and we adopt it to service-oriented geo-processing.

First, user requirements are translated into application requirements. Table 4 exemplifies mapping of user requirements into application level requirements. The table shows that user level *interactivity* requirements are mapped onto *delay* and *reliability* requirements at the application level. The corresponding service levels "high" and "medium" are similarly mapped onto "guaranteed" and "best-effort", and "premier" and "moderate" service levels respectively. What this mapping in essence says is that to achieve, for example, "high" interactivity requires "guaranteed" delay and "premier" reliability. Other user requirements can similarly be mapped.

User::Interactivity	**Application::Delay**	**Application::Reliability**
High	Guaranteed	Premier
Medium	Best-effort	Moderate

Table 4: Interactivity mapped onto delay

The mapping between application level and resource level requirements is a lot more complicated. This is because requirements at the application level refer to a deliverable service and as such, they apply to the entire process of service provision i.e. the resources and the collaborations among them that realize a service. For example the orthophoto service is realized by collaborating mapping, imagery, rectification, and reprojection services, communication links, and geographic datasets (Figure 2). The services, links and datasets are types of resources that are instantiated for consumption in a process to deliver an orthophoto service. Each individual resource typically has its own qualities. Therefore, resource level requirements apply to individual instances of resources and mapping between application level and resource level requirements involves translating process-centric to resource specific QoS requirements.

It is the responsibility of the *geo-service infrastructure* to translate service-centric application requirements into resource-specific requirements. Ideally, this mapping achieves two goals:

- Application requirements are mapped onto aggregate resource requirements
- Aggregate resource requirements are mapped onto individual instances of resources

Table 5 shows a mapping of application level *delay* requirements onto resource level *performance* (a measure of time a resource takes to respond to a request) and *availability* (a measure of the extent to which a resource is available and able to perform required operations) requirements.

Application::Delay	Resource::Performance $\left[P = \sum_{j=1}^{N} p_j\right]$	Resource::Availability $\left[A = \prod_{j=1}^{N} a_j\right]$
Guaranteed	"As-specified"	Guaranteed
Best-effort	Surf-grade	Best-effort

Table 5: Mapping between application and resource level requirements

The *performance (P)* and *reliability (R)* requirements in Table 5 apply to one or more resources as shown by the models beside the requirements. The models define the propagation behaviors of the QoS requirements in a service chain comprising say, "*N*" resources. Guaranteed *delay* means that a service is delivered within specified time constraints hence the performance requirements must be "as specified" i.e. if a *delay* requirement of 30 seconds is specified, then a *performance* of 30 seconds is expected. Also *availability* of resources must be *guaranteed* i.e. resources must be up and running and must execute properly over the entire period they are required to deliver the service. Therefore hard guarantees are provided. On the contrary, for *best-effort delay* requirements no guarantees are provided but every effort is made to deliver the service as soon as possible. Corresponding performance requirements will therefore map to what we call “surf-grade” performance that is the performance experienced when surfing on the Web – a typical best-effort environment.

Application level QoS	Resource level QoS	Propagation model
Delay	Performance (P)	$P = \sum_{j=1}^{N} p_j$ where p_j is performance of resource *j*
	Availability (A)	$A = \prod_{j=1}^{N} a_j$ where a_j is availability of resource *j*
Reliability	Reliability (R)	$R = \prod_{j=1}^{N} r_j$ where r_j is reliability of resource *j*
Cost	Cost (C)	$C = \sum_{j=1}^{N} c_j$ where c_j is cost of using resource *j*
Security	Several security properties can be required e.g. — Data protection — Confidentiality — Authentication — etc	$S = \begin{cases} 1 & if \quad \forall_{(i,j)} \neg \exists_{R_j} \ni q_{(i,j)} = 0 \\ 0 & otherwise \end{cases}$ where $q_{(i,j)} = \begin{cases} 1 & if \quad R_j \quad has \quad P_i \\ 0 & otherwise \end{cases}$ R_j is Resource *j;* P_i is security property *i*
Information Quality (IQ)	— Data quality (DQ) — Computational model quality (CM) — Uncertainty (UN) — Integrity (IT)	$IQ = F(DQ, CM, UN, IT)$

Table 6: Propagation models for resource level QoS

Table 6 summarizes the propagation models for some common resource level QoS requirements including: security; cost; reliability; availability and performance. Information quality (IQ) can be mapped onto data quality, computational model quality, uncertainty (associated either with the data, the model or both), integrity, or any combination of these variables at the resource level. Notably however, the variables are significantly correlated making the models for their propagation behavior and those of their interrelationships non-trivial. The exact propagation models of these variables are outside the scope of this paper suffice it to say that information quality is a function of quality of input data, the quality of the computational model(s), uncertainty (Heuvelink, 1998; Drummond, 1991) and integrity. For security requirements, the basic condition is that all participating resources must support desired security mechanisms or protocols. These are called security properties in the propagation model.

6. CONCLUSIONS

This paper expounded on QoS specification and QoS mapping in the context of service-oriented geo-processing. The orthophoto service was used as a vehicle to demonstrate QoS specification and QoS mapping within the context of a QoS provisioning framework. QoS requirements at different levels of abstraction were defined and mappings between them demonstrated. Application QoS requirements were shown to apply to collections of resources when mapped onto resource level QoS requirements necessitating definition of propagation models for QoS requirements at the resource level. Propagation models for resource level QoS requirements were therefore also defined. Nonetheless, the deployment of effective geo-service infrastructures requires efficient algorithms that make use of the propagation models presented to create an optimal combination of resources that satisfy specified requirement. The algorithms are lacking and present an open research question.

REFERENCES

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

Drummond, J., 1991, Determining and processing quality parameters in geographic information systems (PhD Thesis, University of Newcastle Upon Tyne).

Heuvelink, G.B.M., 1998, Error propagation in environmental modeling with GIS, Londor, Taylor & Francis.

ISO, 2002, Geographic Information: Quality Principles, ISO 19113, Geneve, ISO.

ISO/IEC, 1998, Information Technology – Quality of Service: Framework, ISO/IEC 13236, Geneve, ISO.

Onchaga, R., (2005) (in review), Quality of service provisioning in geographic information service architectures, International Journal of Applied Earth Observation and geo-Information.

Siqueira, F. and Cahill, V., 2000, Quartz: A QoS architecture for open systems, The 20th International Conference on Distributed Computing Systems (ICDCS 2000), Taipei.

Widya, I. and Stap, R.E. and Teunissen, L.J. and Hopman, B.F., 2001, On the End-user QoS-awareness of a Distributed Service Environment, 6th International Conference on Protocols for Multimedia Systems, Enschede.

BIOGRAPHICAL NOTES

Richard is a PhD candidate in the department of Planning and Geo-Information Management (PGM), at The International Institute for Geo-Information Science and Earth Observation, (ITC) Enschede.. He holds an MSc. in Geoinformatics from the ITC, Enchede awarded in 2000 and a BSc. in Surveying obtained from the University of Nairobi, Kenya, in 1994. Richard is a lecturer in the department of Surveying & Mapping , Kenya Polytechnic, Nairobi and has lectured on part-time basis in the Department of Civil Engineering, JKUAT, Nairobi.His research interests include quality of of service (QoS) in distributed GIS, architectures for geographic information services, performance modelling and spatial data infrastructure (SDI).

CONTACTS

Mr. Richard Onchaga
International Institute for Geo-Information Science and Earth Observation (ITC)
P.O Box 6, 7500 AA
Enschede
THE NETHERLANDS
Tel. + 31 053 4874 228
Fax + 31 053 4874 575
Email: onchaga@itc.nl
Web site: www.itc.nl