

A Survey on Service Quality Description

KYRIAKOS KRITIKOS, ICS-FORTH
BARBARA PERNICI, PIERLUIGI PLEBANI, and CINZIA CAPIELLO, Politecnico di Milano
MARCO COMUZZI, Eindhoven University of Technology
SALIMA BENRERNOU, Paris Descartes University
IVONA BRANDIC, Vienna University of Technology
ATTILA KERTÉSZ, MTA-SZTAKI
MICHAEL PARKIN, Tilburg University
MANUEL CARRO, Universidad Politécnica de Madrid and IMDEA Software Institute

Quality of service (QoS) can be a critical element for achieving the business goals of a service provider, for the acceptance of a service by the user, or for guaranteeing service characteristics in a composition of services, where a service is defined as either a software or a software-support (i.e., infrastructural) service which is available on any type of network or electronic channel. The goal of this article is to compare the approaches to QoS description in the literature, where several models and metamodels are included. Consider a large spectrum of models and metamodels to describe service quality, ranging from ontological approaches to define quality measures, metrics, and dimensions, to metamodels enabling the specification of quality-based service requirements and capabilities as well as of SLAs (Service-Level Agreements) and SLA templates for service provisioning. Our survey is performed by inspecting the characteristics of the available approaches to reveal which are the consolidated ones and which are the ones specific to given aspects and to analyze where the need for further research and investigation lies. The approaches here illustrated have been selected based on a systematic review of conference proceedings and journals spanning various research areas in computer science and engineering, including: distributed, information, and telecommunication systems, networks and security, and service-oriented and grid computing.

Categories and Subject Descriptors: H.3.5 [On-line Information Services]: Web-based Services; D.2.1 [Requirements/Specifications]: Languages; C.4 [Performance of Systems]: Modeling Techniques and Performance Attributes; D.2.8 [Metrics]: Performance Measures

General Terms: Documentation, Languages, Management, Performance

Additional Key Words and Phrases: Service, description, provisioning, life-cycle, quality, QoS, model, metamodel, service-level agreement, SLA

The research leading to these results has received funding from the European Community's Seventh Framework Programme FP7/2007-2013 under grant agreement 215483 (S-Cube).

Authors' addresses: K. Kritikos (corresponding author), Institute of Computer Science, FORTH, N. Plastira 100, Vassilika Vouton, GR-700 13 Heraklion, Crete, Greece; email: kritikos@ics.forth.gr; B. Pernici, P. Plebani, and C. Capiello, Politecnico di Milano, Italy; M. Comuzzi, Eindhoven University of Technology, Germany; S. Benrernou, Paris Descartes University, France; I. Brandic, Vienna University of Technology, Austria; A. Kertész, MTA-SZTAKI, Hungary; M. Parkin, Tilburg University, The Netherlands; M. Carro, Universidad Politécnica de Madrid and IMDEA Software Institute, Spain.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© 2013 ACM 0360-0300/2013/10-ART1 \$15.00

DOI: <http://dx.doi.org/10.1145/2522968.2522969>

ACM Reference Format:

Kritikos, K., Pernici, B., Plebani, P., Cappiello, C., Comuzzi, M., Benrernou, S., Brandic, I., Kertész, A., Parkin, M., and Carro, M. 2013. A survey on service quality description. *ACM Comput. Surv.* 46, 1, Article 1 (October 2013), 58 pages.
DOI: <http://dx.doi.org/10.1145/2522968.2522969>

1. INTRODUCTION

A service is an action performed by one entity (the provider) on behalf of another one (the requester) [O’Sullivan et al. 2002]. Through the interaction between these two entities, which is called *service provisioning* and involves various phases, there is a transfer of value from the provider to the requester or recipient. Depending on the service nature and the means or channels available, different service types can be identified. For instance, the drawing of a bank cheque is a physical service that is only available at a bank counter since it requires the requester’s physical presence on a specific place and time to be invoked and delivered. Web services are instead autonomous software systems available over the Internet. This article focuses on software and software-support (i.e., infrastructural) services which are available on any network or electronic channel type. Thus, in this survey the word “service” will have this designated meaning and any other service type will be excluded from the analysis and discussion.

Service orientation has emerged lately as a paradigm facilitating interaction between interoperating systems, but also as a general framework to enable access to IT-based applications, since the benefits of adopting it include interoperability, just-in-time integration, easy and fast deployment, efficient application development, and strong encapsulation [Allen 2006; Georgakopoulos and Papazoglou 2008]. Whilst in the past such interactions were stable and consolidated, several new application environments are now based on access to services with a much shorter time frame and cost, thus responding effectively to ever-changing market conditions, rapid technology improvements, and increased competition and customer needs.

An interesting characteristic of services is that they can be composed of other services (e.g., a transportation service may be composed of land and air transport services). For instance, in the e-business area, services can be selected dynamically and composed in added-value new services, where the composite service components are selected from a number of candidate services offering the appropriate functionality. In many pervasive applications, access to services is based on context characteristics, such as location, environmental parameters, and the like. Utility services, such as the telecom and energy provider ones, also require interoperability of very complex systems to guarantee service delivery. Multimedia and multichannel applications require composing and synchronizing several different services to provide a good user experience. In general, to provide such services, several phases are needed, from the selection of the adequate services to controlling the characteristics of service provisioning, in particular when the providers are not under the direct control of the service user. This composition of autonomously running services requires that the service usage rules specified in agreements are clearly stated and their compliance is verified.

Services can be offered and used across many functional levels following various IT architectures. The functional architecture model considered in this article is a simplified version of the one proposed in Lamanna et al. [2003]. It follows a three-layer architecture. The first layer, the *Application Layer*, contains business or user-oriented applications which may use services to fulfil a part or their whole functionality. The *Service Layer* is the next layer containing services with an electronic interface which are used to build or populate the applications. Various infrastructures, which belong to the *Infrastructure Layer*, host these services and are responsible for managing the

services underlying resources for communication, transactions, security, and so forth, for example, through a platform virtualization environment. Every architecture component can be offered as a service to the same or other component types. For instance, Infrastructure as a Service (IaaS) is the delivery of computing infrastructures as a service which fulfils hosted application or service needs [Dikaiakos et al. 2009].

Since they are intended to be discovered and used by other applications across the Web, services need to be described and understood both in terms of functional capabilities and service quality properties. *Service Quality* is a combination of several qualities or properties (e.g., availability, security, response time) of a service, and can be generally seen as an important factor in distinguishing the success of service providers. The service quality description is the main driver in selecting the best service among a set of functionally equivalent ones. Besides, quality is used to define a contract between a service provider and a service user in order to guarantee that their expectations are met. In addition, such a contract feeds the service management system that is in charge of assessing the proper quality level during the service execution, enforcing it by taking appropriate adaptation actions, such as increasing the underlying service resources, substituting or recomposing the faulty service, and determining which settlement actions apply based on the way the service was executed, such as the final cost or penalties to be paid by the service requester or provider, respectively, and negotiations for contract termination.

In all previous cases, a prerequisite for using service quality is its proper, precise, and rigorous description, covering all possible service life-cycle phases. In this article, the term *quality document* is used to denote the quality description of a given service. This term will be used in a generic way as a description of quality, and all issues related to managing such a document in a specific system architecture will not be addressed. Even not referring to specific architectural solutions, it will be shown that the problem of being able to write quality documents is far from solved. Several approaches have been proposed in the literature, however, there is no commonly agreed way to specify service quality. This article aims at systematically and comparatively reviewing these approaches according to various criteria, including their scope, formality, expressiveness, and applicability. Many important findings are uncovered from the analysis. In addition, areas for further research and investigation are spotted.

As service science and engineering is a new scientific but multidisciplinary field, the authors have performed a systematic review of conference proceedings and journals spanning various related research areas in computer science and engineering, including: distributed, information, and telecommunication systems, networks and security, and service-oriented and grid computing. Only representative research approaches were selected based on the quality of their contribution and their significance and relevance to the scope of this article by also considering the respective publication year, the impact factor of the journal in which they were published (if applicable) and their number of citations. Research approaches published in journals with low impact factor or having a small number of citations were not considered. Depending on the type of the quality document examined, different criteria were exploited to assess its quality, as it will be shown in Section 2, such as the document expressiveness, richness, and formality. In case two or more approaches were assessed to have the same quality, only one was eventually selected having first the highest impact factor (if published in a journal), second the highest number of citations, and third the most recent publication year. In the end, the resulting sample collected is not only representative but also characteristic of the contribution diversity of all relevant approaches in service quality description.

This article is structured as follows. First, in Section 2, a definition of service quality is provided along with an analysis of the service life-cycle and a general classification

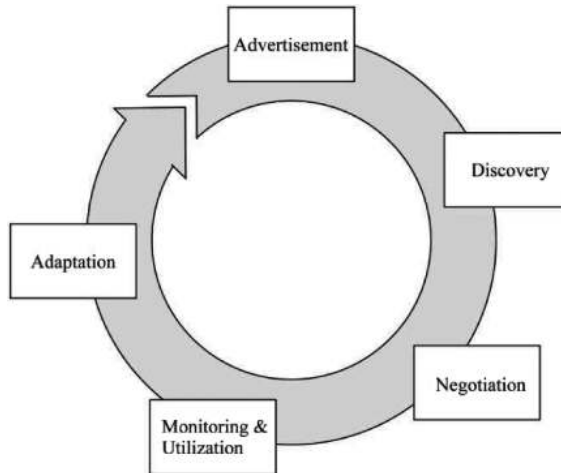


Fig. 1. Service life-cycle.

of quality models is presented, which is then used in the remaining article sections. According to the general classification of quality models presented in Section 2, the state-of-the-art on quality models is analyzed in Section 3, on quality metamodels in Section 4, and on Service-Level Agreements in Section 5. In each section, a set of comparison characteristics are defined and then exploited to compare the approaches. The final section presents interesting topics for further investigation.

2. SERVICE LIFE-CYCLE AND QUALITY

By relying on the definition provided in Kritikos and Plexousakis [2009], which mainly applies to software services, service quality is defined as a set of nonfunctional attributes of those contextual entities that are considered relevant to the service-client interaction, including the service and the client, that bear on the service ability to satisfy stated or implied needs. Service quality can be classified as Quality of Service (QoS) and Quality of Experience (QoE). QoS includes quality attributes that can be objectively measured (like *execution time*) and are typical constituents of Service-Level Agreements (SLAs). QoE includes quality attributes that can be subjectively measured (e.g., *reputation, usability*) and reflects the perception of a service by its users.

The distinguishing feature of service quality with respect to functionality is its dynamicity. In particular, the values of some service quality attributes can vary without impacting the core service function which remains constant most of the time during the service's lifetime. In fact, service quality can play a significant role during several phases of the service life-cycle [Kritikos and Plexousakis 2009]. With reference to the service life-cycle illustrated in Figure 1, each life-cycle activity is shortly analyzed along with the type of quality model it exploits.

- Advertisement*. Requesters and providers publish or exchange quality requests and quality offers, respectively. Such quality documents are called *Quality-Based Service Descriptions* (QSDs).
- Discovery*. The service discovery phase is split into two subphases: service matchmaking and selection. Service matchmaking concerns filtering the advertised services according to the requester's functional and quality requirements. Service selection concerns sorting the matchmaking results according to the requester's preferences.

In result, the requester is presented with an ordered list of services and selects the one that best matches his needs.

- Negotiation*. QSDs are exchanged between service providers and requesters. The possible agreement on quality levels between the parties involved leads to the definition of another quality document, the *Service-Level Agreement (SLA)*.
- Monitoring and Utilization*. The qualities in the SLA are monitored in order to discover customers, and/or providers' violations of its functional and quality terms. Monitoring may also signal potential dangerous situations, which may lead to a violation of the SLA if corrective actions are not timely undertaken.
- Adaptation*. When an SLA is violated, recovery/adaptation actions may be taken. A possible recovery action might require a renegotiation of the SLA or the execution of the matchmaking activity to find an alternative service. It might also happen that an alert is sent to the assessment component of the monitoring activity that continues to execute.

SQMs are descriptions of a taxonomy or concrete list of QoS categories, attributes, metrics, and relationships that connect all of these quality entities. For example, a typical SQM may contain the *Performance* QoS category which includes the QoS attributes of *response time* and *throughput*. As will be shown in Section 3, some proposed SQMs classify quality attributes in terms of relevant scenarios, other SQMs classify them in terms of their dependencies, while other SQMs classify them in terms of compliance to existing standards. Relying on these models means that Service Providers (SPs) and Service Requesters (SRs) have to preliminary select which is the exact set of relevant quality attributes. SQMs provide the concrete semantics of the quality terms that may be used in QSDs and SLAs, that is in other types of quality documents. In this way, all the service life-cycle activities, such as matchmaking and monitoring, are designed around this set of quality attributes. Although the aforesaid procedure assists in producing suitable mechanisms for supporting the service life-cycle activities, the suitability of these mechanisms is specific for the considered scenario.

In addition to QSDs and SLAs, another document type is a Service Quality Model (SQM). SQMs have been used to describe concrete quality properties which can be exploited by other quality document types to express service quality capabilities/requirements or service levels. QSDs have been used to express service quality capabilities and requirements as a set of constraints on quality attributes and metrics. Such descriptions contain all the appropriate information for supporting the service matchmaking and negotiation activities. Finally, SLAs have been used not only to express the service levels in which a service can execute but also other information that is suitable for supporting the service provisioning activities. As a service-level description of an SLA is actually a QSD (i.e., a set of service quality constraints), SLAs are actually built on top of the other document types.

QSDs are often associated with a *validity period* or *expiration time* which signifies when they become outdated. Depending on which party is producing them, QSDs can be separated into *Service Quality Offers* (produced by an SP) and *Service Quality Requests* (produced by an SR). Service Quality Requests are further separated into *Service Quality Requirements* and *Service Selection Models*. The latter denote the significance of each quality attribute or metric to the SR by associating it with a specific weight and are used for ranking Service Descriptions (SDs). Both Service Quality Offers and Requirements are expressed as a set of *quality constraints*. A quality constraint usually contains a comparison operator that is used to compare a quality metric or attribute with a value. Sometimes, a quality constraint may also contain the unit of the compared value. Thus, QSDs describe all the appropriate information that is required for matchmaking and negotiating service quality. In this way, they are used in the respective

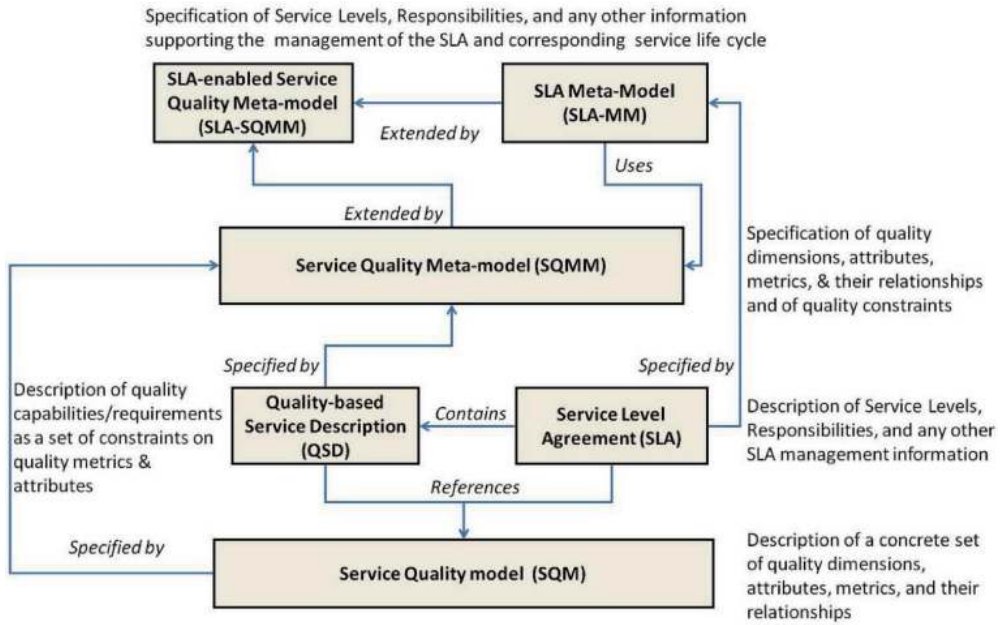


Fig. 2. The main types of quality documents and metamodels and their interrelationships.

service life-cycle activities. However, they are not used further in the service life-cycle, as they do not contain all the appropriate information that is required for supporting the rest of the service life-cycle activities, which mainly concern the service provisioning.

To this end, SLAs have been introduced to close this gap. An SLA is an important aspect of a contract for IT services that includes the set of QoS guarantees and the obligations of the various parties (see definition in Keller and Ludwig [2003]). QoS guarantees are widely known as Service-Level Objectives (SLOs) and are expressed as conditions on one or more QoS metrics, thus indicating the metrics, allowed values. A set of SLOs constitutes a specific Service Level (SL). There can be different SLs defined in an SLA, expressing the different modes in which a service may be executed in different time periods, or degradation/upgrade levels if the agreed SL is violated/surpassed. The party obligations are usually expressed as action guarantees (e.g., compensation, recovery, or management actions) to be performed when a given precondition is met (e.g., when a violation occurs). Other important SLA components are the organizational ones which correspond to information concerning service monitoring and reporting.

As can be seen from the preceding analysis, SLAs contain more information than QSDs in terms of supporting the service provisioning activity. Moreover, there is no uniform and common quality document to be used across all the activities. This is a major drawback that requires time, as document transformations should take place from one format to the other, and reduces the automation degree of the activities.

In order to automate as much as possible the aforesaid activities, a clear and formal description of QoS is required. Moreover, Service Providers (SPs) and Service Requesters (SRs) should agree on the same language for expressing their quality documents. In this way, all the mechanisms used for supporting the service life-cycle can be properly enacted. Nowadays, in the literature many metamodels and languages for describing service quality have been proposed, which can be distinguished in two main types: *Service Quality MetaModels* (SQMMs) and *SLA MetaModels* (SLA-MMs). Figure 2 represents the types of quality documents and metamodels and the various relationships involved between these entities.

SQMMs provide the means for describing QoS in a more general and extensible way than SQMs. Actually, an SQMM is a conceptualization of the appropriate quality concepts and their relationships that can be used to capture and describe an SQM. For example, a typical SQMM will contain the concepts of *QoS category*, *QoS attribute*, and *QoS metric* and the relationships *contains* (from QoS categories to attributes) and *measuredBy* (from QoS attributes to metrics). So, an SQMM can describe many different SQMs, where the number of those SQMs and their actual difference mainly depends on the richness of the SQMM. In addition, SQMMs are used to specify *QSDs*, which are usually described by a set of constraints on some QoS attributes and metrics. Thus, existing SQMMs can be compared according to their expressiveness, as will be shown in Section 4. On one hand, by adopting an SQMM, the mechanisms that use quality documents become more general than those which adopt a specific SQM. Indeed, in this case, these mechanisms can be designed regardless of a specific quality attribute set. In case SRs and SPs change the relevant quality attribute set, the mechanisms remain the same. On the other hand, due to the intrinsic subjectiveness and complexity of quality, the existing SQMMs are not able to capture all the possible features of quality attributes and rely on a common understanding of the interacting parties about the concepts defined in the SQMM.

As will be discussed in Section 5, SLA-MMs allow the definition of *SLAs* and *SLA Templates* between the interacting parties. Since the agreement terms include Service-Level Objectives (SLOs), which denote constraints on quality attributes or metrics listed in an SQM, and both SQMs and constraints may be defined by an SQMM, the existence of the three following cases is highlighted: (a) there is a specific SQMM type, called *SLA-enabled SQMM* (SLA-SQMM), that can define SLA specifications; (b) SLA-MMs may use one or more SQMMs to define and reference quality attributes and even specify SLOs; (c) SLA-MMs may reference the contents of one or more SQMs. Various SLA description capabilities are considered, when comparing existing SLA languages, which concern the definition of the contract terms and various other information that may be used to support the service life-cycle activities.

In general, metamodels are used to define a language's abstract syntax. Then, different concrete syntaxes may exist for the same language, which are based on its abstract syntax. Therefore, a metamodel drives the design of a language. Thus, a language has one and only one underlying metamodel, whilst one metamodel may be used for the design of many languages. However, in practice, there is usually a one-to-one correspondence between a language and its underlying metamodel, as different languages of the same domain are designed by different modelers which have a different conceptualization of that domain. Indeed, to the best of our knowledge, *Service Quality Specification Languages* (SQSLs) and *SLA Languages* have an one-to-one correspondence with their underlying SQMMs and SLA-MMs, respectively. Thus, these corresponding terms will be used interchangeably in the remaining sections of this survey.

3. SERVICE QUALITY MODELS

3.1. Background

As described in Section 2, an SQM focuses on the analysis of the set of quality attributes that are considered relevant in service applications. If services are considered as standalone software modules, then their quality can be determined by the attributes that traditionally characterize software quality and, thus, by the attributes defined in the ISO 9126 model [ISO/IEC 2001]. However, this model must be adapted to capture the peculiarities deriving from the service intended use, that is, their composition to build service-based applications by removing the nonapplicable attributes and their categories (i.e., most of the internal attributes) and refining or specializing other attributes (i.e., the rest of internal and all external attributes).

Services in the advertisement and utilization phase are considered as black boxes that expose their external attributes to the audience of service-based developers. From this perspective, all external quality attributes proposed for traditional software are applicable, such as *privacy*, *security*, *performance*, and *reliability*. However, some internal quality attributes are not applicable (e.g., *analysability*, *changeability*) since they require the analysis of the software code that is hidden in the service-oriented programming or consider *portability* aspects which are de-facto covered through the use of the standard service technology. On the other hand, some internal software product characteristics can be useful at least during service design as they influence the way a composite service is built or executed. For instance, in case of the *stability* attribute, a composite service developer would select more stable service components than others to build at design or runtime a stable composite service which can be used many times without being updated.

Thus, it is easy to understand that the ISO 9126 quality model is not adequate for representing service quality. In addition, it applies only to software services and not to other service types, such as infrastructural services. For this reason, different contributions can be found in the literature, which propose various SQMs. The structure of these SQMs is based on the use of taxonomies in which categories, related to different analyzed aspects, are defined. Each category contains a set of attributes that are entities which can be verified or measured in the service. Most of the models associate each attribute with a definition and, in some cases, also provide the related metric and assessment formulae. The latter information is needed only for measurable quality attributes [Kritikos and Plexousakis 2009].

The most referenced QoS categories of SQMs are the *Performance*, *Security*, *Dependability*, and *Configuration* categories that usually contain specific quality attributes, as will be shown in Section 3.3. As there is a differentiation on how SQMs categorize the rest of the service quality attributes, it was decided not to evaluate SQMs in terms of specific categories but generally on the extensiveness of their quality attribute categorization. Moreover, SQMs were evaluated also on the level of detail in their categorization in order to inspect the SQM richness. Section 3.3 evaluates SQMs in terms of the quality attributes they contain in order to distinguish which attributes are the most common ones and to provide a proof for the correct evaluation of the SQM richness.

3.2. Methodology

The most significant SQMs proposed for services and their applications have been collected according to the selection process designated in Section 1. Only generic SQMs were considered and not specialized models proposed for security, data, and network aspects. This is because in the analysis of each security and data quality model the service-level and service type attributes are not considered, since the respective classifications have been often defined for generic applications and not for service-based ones.

A set of comparison criteria have been selected in order to analyze and compare the different SQM approaches according to the type and value of the information that they contain. In particular, a set of particular aspects of the information that such an SQM must specify were identified by considering specific requirements originating from the service life-cycle presented in Section 2. These include the consideration of both the service provider and requester view for service discovery and negotiation and the association of a quality attribute with assessment guidelines for enabling its assessment, as well as quality-based specific criteria, such as the extensiveness of the SQM's quality categorization and the consideration of both domain-independent and domain-dependent attributes. Some of these criteria reflect the directions for further research with respect to service quality models indicated in Seth et al. [2005], such

Table I. Summary of the Comparison Criteria of Service Quality Models

Criterion	Summary
Quality Categorization Extensiveness	How many quality categories does the SQM have?
Level of Detail in Quality Categories	Does each quality category contain a sufficient number of quality attributes?
Containment of Domain-Independent and Domain-Dependent Quality Attributes	Does the SQM contain both domain-independent and domain-dependent quality attributes?
Consideration of Service Provider and Requester Views	Are the contained quality attributes relevant to the SP, SR, or both?
QoS and QoE Consideration	Does the SQM contain both quality attributes that are measured objectively and those that are measured subjectively?
Atomic and Composite Quality Attribute Inclusion	Are there any relationships expressed between atomic and composite QoS attributes?
Types of Dependencies	Is there any type of inter-attribute dependencies expressed?
Layer Designation	To which service layer(s) does a QoS attribute refer to?
Association with Assessment Guidelines	Are there any assessment guidelines associated to a QoS attribute?
Metric Identification	Is there any metric associated to a QoS attribute?

as measurement issues (quantification and measurement of QoS and linkage of QoS with business performance—criteria Layer Designation, Association with Assessment Guidelines, Metric Identification, Types of Dependencies), relation between various attributes of a service (criterion Types of Dependencies), and role of technology such as IT (criteria—Layer Designation and Types of Dependencies). It must be noted that the criteria considered do not assess which exact modelling constructs are used by a SQM to specify appropriate quality terms and their relationships as this is the subject of research for the SQMM evaluation in Section 4.

A summary of the comparison criteria considered is shown in Table I, while their thorough presentation is provided later on in this section. The evaluation results of the examined SQMs according to the comparison criteria are shown in Table II. In this table, the SQMs are sorted in ascending chronological publication order. In this way, interesting conclusions concerning trends in SQM modelling can be drawn.

In the remainder of this section, each criterion along with its evaluation results is presented in separate subsections. In the end, a global analysis of the SQMs across all criteria is given and the most frequent quality categories and attributes in the considered SQMs are distinguished.

3.2.1. Quality Categorization Extensiveness. Quality attributes are usually categorized into quality categories. This categorization is required as the set of quality attributes is often large and can be used to separate quality according to several aspects. In this way, the higher the number of categories, the higher the comprehensiveness of the approach. Categories also improve the model's readability and allow users to better explore the model and optimize the search activity. Categories that are defined in many SQMs are *Performance*, *Security*, *Dependability*, and *Configuration Management*. A flat model would contain only one quality category while an extensive one can contain up to 9 or 10 categories. Grades of extensiveness are defined as follows: *flat* (1 category), *fair* (2–4 categories), *good* (5–7) categories, and *extensive* (8–10) categories. While it is quite obvious that a flat SQM would contain just one quality category, the other SQM space partitions were created in a uniform way by observing that the highest number

Table II. Comparison of Research Approaches Providing Service Quality Models

Research Approach	Extensiveness	Detail Level	Domain Dep/ncy	SP/SR View	QoS & QoE	Composite & Atomic	Dep/ncy Types	Layer	Metric	Ass/ent Guid/es
[Sabata et al. 1997]	fair	fair	domain ind.	SP	QoS	both	no	service	partial	fair
[Ran 2003]	good	fair	domain ind.	SP	QoS	both	no	all	partial	fair
[Colombo et al. 2005]	flat	high	domain ind.	both	both	atomic	no	serv. & appl.	complete	fair
[The OASIS Group 2005]	fair	fair	domain ind.	both	QoS	both	quant.	serv. & appl.	complete	good
[Cappiello 2006]	extensive	high	domain ind.	both	both	both	no	all	partial	none
[Truong et al. 2006]	good	fair	domain ind.	SP	QoS	both	quant.	serv. & infr.	partial	fair
[Brandic et al. 2006]	good	fair	domain ind.	SP	QoS	both	quant.	serv. & infr.	partial	fair
[Sakellariou and Yarmolenko 2008]	good	good	both	both	QoS	both	no	serv. & infr.	partial	fair
[Cappiello et al. 2008]	extensive	high	both	both	both	both	quant.	all	none	none
[Frutos et al. 2009]	good	high	domain ind.	SP	QoS	both	quant.	all	partial	none
[Nessi Open Framework 2009]	fair	fair	domain ind.	both	both	both	quant.	serv.	none	none
[Kritikos and Plexousakis 2009]	extensive	high	both	both	both	both	no	all	partial	fair
[Mabrouk et al. 2009]	extensive	good	both	both	both	both	quant.	all	partial	fair

of quality categories is 10 across all SQMs and splitting the [2–10] space of the number of quality categories in three (additional) equally-sized partitions.

The evaluation results are presented in the second column of Table II. The majority of SQMs use categories to classify the attributes and improve the model’s understandability. Moreover, there is a balance between all the SQM partitions created according to this criterion (apart from the *flat* partition which has only one approach) with the partition corresponding to good extensiveness to have a very small precedence. Finally, by observing the last nine values of the extensiveness criterion column, we can see that SQMs are improving according to this aspect (with the exception of Nessi Open Framework [2009]).

3.2.2. Level of Detail in Quality Categories. This criterion is used to determine if each category contains a sufficient quality attribute number. The greater the number of quality attributes in a category, the higher the probability that the most significant QoS attributes will be included within it. So when a quality category contains less than three attributes, its *level of detail* is: *low*. When it contains three to five attributes, it is *good*, while when it contains more than five it is *high*. Again, the splitting of the quality category space according to the number of the quality attributes contained was performed in a uniform way based on the lowest and greatest number of quality attributes across all quality categories of all SQMs. In this way, we can build on this categorization to express the *level of detail* of the whole SQM. To this end, the following

levels have been defined by considering all meaningful combinations of categorizations of an SQM's quality categories: *very low* (every category has low *level of detail*), *low* (some categories have low and other have good *level of detail*), *fair* (some categories have low and other have good or high *level of detail* – but at least one quality category has high *level of detail*), *good* (all categories have good *level of detail*), *very good* (some categories have good and other have high *level of detail*), *high* (most categories have high *level of detail*).

The evaluation results are presented in the third column of Table II. Almost half of the SQMs have a fair level of detail which means that some of their categories contain less than three quality attributes. For the rest of the approaches, SQMs that present high level of detail are more than those having good. Finally, note that approaches with a level of detail greater than fair have appeared after 2005, meaning that SQMs have increased their level of detail over time.

3.2.3. Containment of Domain-Independent and Domain-Dependent Quality Attributes. Domain-independent quality attributes are general/technical attributes that can characterize all services in any possible application domain. For example, it is difficult to find an SQM that does not contain the *response time* and *availability* (domain-independent) attributes. On the other hand, domain-dependent quality attributes characterize services (or their parts) of one or more but not any application domain. For example, data-related attributes like *validity* and *timeliness* characterize the input or output data of services that manipulate data. So the evaluation of this criterion for a specific SQM would be: *domain ind.* (only domain independent), *domain dep.* (only domain dependent), and *both*.

The evaluation results are presented in the fourth column of Table II. All the approaches tend to be general and not domain specific. In fact, all the SQMs propose domain-independent criteria that can be generally used in every context in which nonfunctional properties are considered. Only four approaches enumerate also some domain-specific attributes of a particular domain [Sakellariou and Yarmolenko 2008; Kritikos and Plexousakis 2009; Mabrouk et al. 2009] or some quality attributes used in some domains according to a particular context [Cappiello et al. 2008]. The latter four approaches have been proposed recently.

3.2.4. Consideration of Service Provider and Requester Views. This criterion considers the relevance degree of the different quality attributes with respect to the SP and SR. Some quality attributes are particularly relevant for SPs, such as *availability* and *response time*. Other quality attributes are important to SRs, such as *response time* and *usability*. Thus, some quality attributes are both important to SPs and SRs, so they belong to their *common* view, while other attributes are important only for one of these parties, so they are either *SP specific* or *SR specific*. If the SQM captures both views, it should provide a quality attribute set that can be used to express both SP capabilities and SR requirements. Thus, this criterion's evaluation for a specific SQM would be: *SP* (if the SQM contains *SP-specific* and *common* quality attributes), *SR* (if it contains *SR-specific* and *common* attributes), and *both* (if it contains *common*, *SP-specific*, and *SR-specific* attributes).

The evaluation results are presented in the fifth column of Table II. More than half of the SQMs consider quality attributes that correspond to both the provider's and requester's view. So, researchers have understood the need of expressing both views in an SQM. Moreover, the evaluation results of this criterion are in accordance with those of the previous one. First, domain independence signifies that the selected quality attributes are more specific to the service and independent of its usage in specific applications, so at least the provider view is covered. This is because providers are more concerned about having their services used across many application domains so they

focus more on those quality attributes that are generic and tend to distinguish their services from other services independently of the application domain. Second, domain dependence signifies that the selected attributes are specific to an application and its usage, so they have an impact on the user's expectations. Thus, domain-dependent attributes tend to cover the requester view. Indeed, by cross-checking the results of this and the previous criterion (domain independence), we can observe that when only domain-independent attributes are contained in an SQM, then at least the SP's view is captured. In addition, when an SQM also captures domain-dependent attributes, then also the requester's view is captured. Finally, it can be observed that while in the past most SQMs were capturing only the SP view, this situation has changed recently as the more recent SQMs (from 2006 and on) tend to cover both views.

3.2.5. QoS and QoE Consideration. QoE attributes certainly reflect the requester view. However, two main questions arise that may be addressed by the evaluation of this criterion and the observation of the evaluation results of the previous criterion: (a) are all requester-view quality attributes QoE or not?, and (b) what is the situation with the provider-view quality attributes, in other words do the provider-view attributes contain QoS, QoE, or both attribute types? Thus, these two questions actually concern which views do the QoS attributes reflect. Obviously, both sets of attributes are important and should be represented by an SQM. So, the evaluation of this criterion for a specific SQM would be: *QoS* (only QoS), *QoE* (only QoE), and *both*.

The evaluation results are presented in the sixth column of Table II. It is noticeable that there is a balance between the SQMs that contain only QoS attributes with those that contain also QoE attributes. Moreover, the more recent SQMs tend to cover both attribute types. This means that SQM modellers have understood the importance of modelling both QoS and QoE attributes.

By inspecting this criterion's results and those of the previous one, some other interesting facts can be inferred. First, when an SQM covers only the SP view, then it covers only QoS attributes. This means that an SP considers the QoS attributes as more important. This is quite reasonable as the SP-view quality attributes are in their majority domain-independent attributes that should be measured objectively in order to be able to meaningfully compare services across all application domains. Second, when an SQM additionally covers the SR view, then it covers either QoS or both QoS and QoE attributes. This signifies that the SR view corresponds to both attribute types. On one hand, the SR-view domain-independent attributes will tend to be QoE attributes because these attributes can be assessed differently from users across the various application domains, as in each domain the usage and the requirements from a service are different with respect to the other domains. On the other hand, SR-view domain-dependent attributes will tend to be QoS attributes because in a specific application domain the requirements and the service usage are specific or vary in a specific way according to the user expectations, while the user's domain expertise is high. In this way, the SR-view domain-dependent attributes will tend to be measured with well-established domain metrics or be assessed in the same objective way by users that have similar expectations.

3.2.6. Atomic and Composite Quality Attribute Inclusion. Some attributes are *composite* and can be computed by evaluating other attribute values. For example, *response time* (the parent) is a composite attribute since it can be assessed by evaluating *latency* and *network delay* (its children). Other attributes, like *authentication*, are *atomic* as they do not rely on any attribute. The composability aspect is important as it can be used later by metamodels to capture this inherent relationship type between the parent and child quality attribute. Moreover, as this relationship can be associated with specific quantitative (composite attribute derivation formulas) or qualitative (expressing the

correlation between the parent and the child attribute) dependencies, it can provide important input to the service monitoring process. Thus, this criterion's evaluation for an SQM would be: *atomic* (only atomic attributes are included), *composite* (only composite attributes are included), and *both* (both attribute types are included along with a connecting relationship).

The seventh column of Table II presents the evaluation results. Most of the SQMs contain both atomic and composite quality attributes along with a connecting relationship. In addition, all SQMs proposed after 2005 contain both attribute types, demonstrating an improving trend with respect to this aspect. Finally, by inspecting the evaluation of the SQM of Colombo et al. [2005] according to this and the first two criteria, it can be observed that this SQM contains only atomic attributes, while its category extensiveness is flat and level of detail is fair. Such an SQM is built and structured in this way for supporting fact computation (i.e., service matchmaking and selection) algorithms which require that no interattribute dependencies exist but only QoS attributes that can be immediately measured by specific metrics without introducing higher measurement levels.

3.2.7. Types of Dependencies. There can be two dependency types between quality attributes: *quantitative* and *qualitative*. The former are expressed by mathematical formulas or constraints, while the latter are expressed symbolically or descriptively. A qualitative dependency example is that *availability* and *reliability* have "a positive correlation", that is, an increase of the one's value causes an increase to the other's value. Both dependency types can be exploited for service monitoring. For example, the monitoring process can inspect if the monitoring facts conform to the correlations expressed in qualitative dependencies. In addition, the dependencies between quality attributes may unveil the correlation between quality attributes across the three service layers and, thus, provide insights on how the performance of an application or business process is influenced by the performance of the (external) services that it exploits as well as by the performance of the infrastructure that provides support to its execution. So it must be inspected if concrete dependencies between quality attributes exist without considering which modelling constructs are used, as SQMs and not SQMMs are evaluated. This criterion's evaluation for a specific SQM would be: *no* (no dependencies are expressed), *quant.* (only quantitative dependencies are expressed), *qual.* (only qualitative dependencies are expressed), and *both* (both dependency types are expressed).

The evaluation results are presented in the eighth column of Table II. Only quantitative dependencies are addressed by the SQMs. In addition, all of these quantitative dependencies concern composite attribute derivation formulas. Finally, by comparing the evaluation results with those of the previous criterion, it can be inferred that the majority of the SQMs that specify both composite and atomic quality attributes also specify the way the composite attributes are produced from the atomic ones.

3.2.8. Layer Designation. The reference model introduced in Section 1 consists of three layers: *application*, *service*, and *infrastructure*. The *application* and *service* layers usually have a set of identical attributes, where the application-layer attributes are produced from the respective service-layer ones. Moreover, some business- or user-oriented attributes are associated to the application layer. The infrastructure layer usually contains a completely different quality attribute set with respect to the other two layers. So, as it is important to clarify to which service layers an SQM refers, this criterion's evaluation for a specific SQM would be: *service*, *serv. and appl.* (service and application), *infr.* (infrastructure), *serv. and infr.* (service and infrastructure), and *all* (all layers). We have considered all meaningful layer combinations for categorizing the SQMs, where the term "meaningful" relies on the fact that as we are dealing with the quality of a

service, at least the service layer must be captured by an SQM through defining the respective quality attributes of this layer.

The evaluation results are presented in the ninth column of Table II. As all SQMs focus on the components of the service layer, that is, software services, they unsurprisingly contain service-layer quality attributes. Almost half of the SQMs also contain either infrastructure-layer or application-layer attributes, while the same result applies for SQMs which have adopted a holistic approach. Thus, apart from the initial SQM that contained solely service-layer quality attributes, the researchers in this field have quickly understood the need to include quality attributes from other layers so as to be able to characterize the quality of all types of services.

By comparing the results of this criterion with those of the previous one, it can be inferred that there is a trend in proposing SQMs that not only cover all service layers but also expose the quantitative dependencies between the service attributes. However, it should be noted that most of the SQMs do not cover dependencies between quality attributes at different service layers but only at the same layer. In this way, first the lack of qualitative dependencies prevents monitoring from deriving novel facts or validating the existing ones. Second, the way the quality at one layer is measured and influenced by the quality at a different layer is not indicated, thus preventing the service monitoring and assessment activities to correlate the monitored information, discover interesting patterns, and find the root cause of a specific service fault that has occurred with respect to the violation of a specific SLO of an SLA. Indeed, by also considering the state-of-the-art work in service monitoring, it can be observed that cross-layer service dependency models are scarcely exploited and usually the monitoring components of an application are deployed at specific layers and are usually independent from each other. As such, they cannot easily detect which is the exact root cause of a service-level degradation and take the most suitable adaptation action or strategy. Moreover, they may detect different forms of the same problem and initiate independent adaptation actions which may have strong interference or cancel the effects of each other, and may potentially not solve the problem in its entirety [Kazhamiakin et al. 2009].

3.2.9. Metric Identification. This criterion is used to inspect if SQMs contain quality metrics used to measure quality attributes. Metrics are entities that encapsulate all appropriate measurement details of an attribute such as measurement values, units, formulas, and schedules. However, a metric's measurement formula or assessment algorithm is used as a separate comparison criterion from the current one because some SQMs provide the name of a metric and some of its details, when they associate it to an attribute, but not an assessment guideline or algorithm for it. So it must be assessed first which SQMs associate specific metrics to attributes and then the existence or not of assessment guidelines for the incorporated metrics. Thus, this criterion's evaluation for each SQM will take the following discrete values: *none* if no metrics are provided in the SQM, *partial* if the SQM associates metrics to a subset of the quality attributes, and *complete* when all attributes are associated with at least one corresponding metric.

The evaluation results are presented in the tenth column of Table II. More than half of the SQMs associate metrics to a subset of the contained quality attributes, while only two SQMs are complete in this aspect. The reason of having many *partial* SQMs can be twofold: (a) some quality attributes are not measurable or are difficult to measure (e.g., QoE attributes), and (b) there was a decision of associating metrics only to the most popular or widely used (in QSDs and SLAs) attributes. Finally, two SQMs do not associate metrics to attributes, which significantly limits their usage in specific service quality description and matchmaking scenarios.

3.2.10. Association with Assessment Guidelines. An SQM may provide algorithms to assess the quality attributes that it defines. This criterion's evaluation would be: *none* if guidelines are not provided, *fair* if simple assessment rules are provided, and *good* if the authors specify precise assessment algorithms.

The evaluation results are presented in the eleventh column of Table II. More than half SQMs just specify some guidelines for each metric, while one third of the considered SQMs do not provide metric assessment guidelines. Thus, as the latter SQMs cannot be used in the monitoring and assessment service life-cycle activity, their use is limited mostly in the advertisement and matchmaking activities. Only one approach provides precise assessment algorithms for all defined metrics and associates metrics to all its contained attributes. So, this approach is suitable for annotating QSDs and SLAs which can be used across all service life-cycle activities. However, it can be used in specific scenarios as it contains a small set of domain-independent quality attributes and not domain-dependent ones.

By correlating the evaluation results of this criterion with those of the previous one, it can be inferred that SQMs, which associate metrics to some or all of their quality attributes, do provide for them simple assessment rules that could be further used to create precise assessment algorithms. If this inferred result is combined with that for the Types of Dependencies and Service-Layer criteria, it can be observed that most SQMs do not provide the appropriate information for the proper service monitoring and assessment of a service's quality attributes. This substantially limits their use and creates the need for a novel SQM that not only extends the best SQM according to these aspects but also provides quantitative and qualitative dependencies between quality attributes of not only the same but of also different service layers.

3.3. Overall Analysis

Based on the preceding analysis, no SQM can be distinguished as the best according to its evaluation on all the considered criteria. On one hand, by considering the first six criteria plus the Layer Designation one which relate to an SQM's extensiveness, structure, and generality, four approaches can be distinguished as the most complete [Cappiello 2006; Cappiello et al. 2008; Kritikos and Plexousakis 2009; Mabrouk et al. 2009]. On the other hand, by considering the last two criteria which relate to an SQM's attribute assessment and applicability, the approach in Colombo et al. [2005] can be distinguished. Moreover, all approaches have the worst behavior with respect to the Types of Dependencies criterion. Thus, a new SQM is needed combining the characteristics of the best approaches in the aforesaid two criteria partitions and describing all the possible but realistic interattribute dependencies across the same and different service layers. Such an SQM could be specified via the ideal metamodel illustrated in Figure 3, where the concepts and associations marked with grey show the common conceptual elements among all SQMs.

Apart from the criteria-based analysis, the frequency of the service quality categories and attributes across all SQMs was assessed in order to distinguish the most frequent ones. Table III shows which quality attributes have been defined by which SQM(s).

Although several different QoS attributes and categories can be found in the various proposals, it is possible to single out the most frequent and consider them as the "basic" and most important QoS attributes and categories, respectively. This is because all other attributes either capture secondary features or are more context dependent (i.e., very specific), while they appear very scarcely in the proposed SQMs. As can be seen from Table III, *response time*, *latency*, and *throughput* are the attributes that mostly represent the *performance* category, which is present in most SQMs. Another important category is *security* which has three attributes, namely *authentication*, *authorization*,

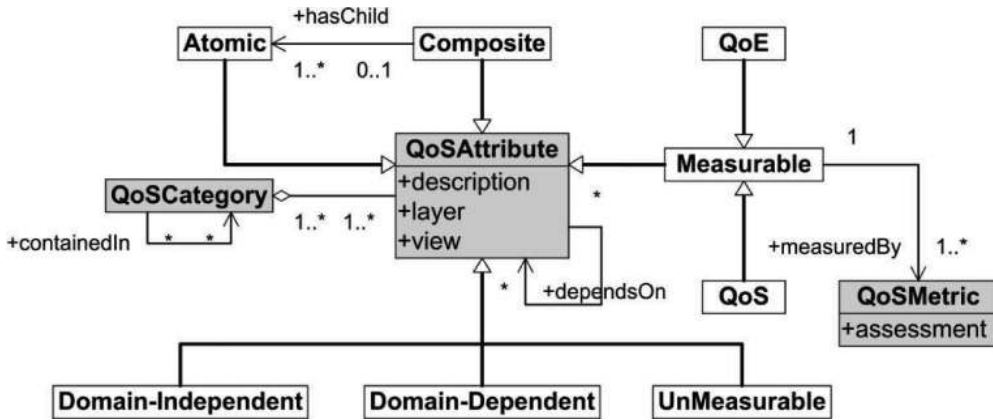


Fig. 3. Ideal metamodel for SQMs.

and *nonrepudiation*, that are steadily present in the SQMs. *Availability*, *accuracy*, and *reliability* are the remaining three most important attributes.

Data quality is a multidimensional concept that defines the suitability of the used data for the application involved. In the literature, there exist several contributions about data quality attributes and taxonomies/SQMs (e.g., Redman [1997], Strong et al. [1997], and Batini et al. [2009]). The most important and representative data quality attributes (according to the data quality research) are *accuracy*, *completeness*, *consistency*, and *timeliness*. Since the service output is mostly composed of information, data quality can be considered as a part of the service QoS and it can thoroughly drive the analysis of the required input and provided output. Thus, data quality attributes could be easily applied to the service world in order to check the correctness of data, the existence of missing or contradictory values, and the updateness of the information provided. However, data quality aspects are scarcely considered. The only data quality aspect that is mostly considered is *correctness*.

Some SQMs take into account particular network aspects. In these SQMs there is usually a *Network* quality category which mainly contains the four most frequent network attributes, namely, *bandwidth*, *network delay*, *jitter*, and *packet loss*.

4. SERVICE QUALITY METAMODELS

4.1. Background

SQMMs have been mainly used to describe the service quality capabilities or requirements of an SP or SR, respectively. Thus, apart from their ability to describe SQMs, SQMMs can specify *QSDs*. As service description is a prerequisite for service discovery, the content of SQMMs has been used for Quality-Based Service Matchmaking (QBSM) and service selection in service registries. QBSM is a process executed after Functional Service Discovery (FSD) to further filter out a registry’s Service Descriptions (SDs) based on their service quality capabilities with respect to the service quality requirements of an SR. It must be noted that SDs specify both the service functional and quality capabilities. The results of the QBSM process may be ranked, if needed, by executing the service selection process.

As analyzed in Section 2, some SQMMs can be also considered as SLA-MMs because they can describe SLAs. These SQMMs are called *SLA-enabled SQMMs* (SLA-SQMMs). The corresponding languages of this SQMM type include QML [Frølund and Koistinen 1998], WSLA [Keller and Ludwig 2003], WSOL [Tosic et al. 2003b], and

Table III. Attributes Defined in Service Quality Approaches

Service Quality Category	Service Quality Attribute	Approach Reference											
		Sabata et al. 1997	Ran 2003	Colombo et al. 2005	The OASIS Group 2005	Cappiello 2006	Truong et al. 2006	Brandic et al. 2006	Sakellariou and Yarmolenko 2008	Cappiello et al. 2008	Frutos et al. 2009	Nessi Open Framework 2009	Kritikos and Plexousakis 2009
Performance	Response Time		X	X	X	X	X		X	X	X	X	X
	Processing Time					X	X	X	X	X		X	X
	Latency		X	X		X	X		X	X		X	X
	Timeliness	X		X		X			X				
	Precision	X											
—	Throughput		X	X	X	X			X	X	X	X	X
	Availability		X	X	X	X	X		X	X	X	X	X
	Accessibility				X	X	X		X	X		X	X
	Accuracy	X		X		X	X		X	X		X	X
	Reliability		X	X		X	X		X	X	X	X	X
	Capacity		X				X		X	X		X	X
	Believability						X						
	Maintainability			X			X					X	X
	Relative Importance	X											
	Complexity			X									
	Customer Service			X									
	Depandability			X		X			X		X	X	X
	Stability		X	X	X				X			X	X
	Trust			X		X							
	Understandability			X					X				
	Integrability											X	
	Interoperability					X						X	
	Resource Efficiency											X	
	Reusability											X	
	Scalability		X			X			X		X	X	X
Cost	X	X			X	X	X	X	X	X		X	X
Security	Authentication		X		X	X	X		X	X	X	X	X
	Authorization		X		X		X		X	X	X	X	X
	Level	X					X					X	X
	Integrity		X		X		X		X			X	X
	Confidentiality	X	X		X	X	X		X	X		X	X
	Accountability		X				X		X	X		X	X
	Traceability		X		X				X	X		X	X
	Non-repudation		X		X	X			X	X	X	X	X
	Data Encryption								X	X	X	X	X
Configuration	Isolation										X		
	Virtual Organization					X	X	X					
	Location						X	X					
	Level of Service	X					X		X	X			
	Service Version						X						
Data	Supported Standard		X	X	X		X		X			X	X
	Maturity/Age			X		X			X				
	Timeliness					X			X			X	
	Completeness					X			X	X		X	

Table IV. Summary of the Comparison Criteria of Service Quality Specification Languages

Criterion	Summary
Formalism	The language's expression formalism
Quality Model	Expressiveness in defining SQMs
Metric Model	Expressiveness in defining metric models
Constraint Model	Expressiveness in defining service quality constraints
Complexity	The complexity of producing SQMs and QSDs
Service Description Separation	Separation of functional and quality-based SDs
Service Description Refinement	Refinement/reuse of QSDs
Service Description Granularity	The ability to define quality constraints for the various service components
Symmetric but Separate QSDs	QSDs should be defined for both SPs and SRs in the same way but separately
Class of Service	The ability to produce different QSDs for the same service
Connection	Connection of a language's QSDs with functional SDs of a specific language
Quality Matching	The way QSDs of SPs and SRs should be matched
Framework Support	In which type of frameworks is the language used

SWAPS [Oldham et al. 2006]). SLA metamodels that do not define quality attributes and the corresponding service quality capabilities are not considered in this section (e.g., WS-Agreement [WS-AGREEMENT 2003]). Section 5 will analyze all kinds of SLA languages with the appropriate SLA-based criteria.

Some security aspects like *trust* and *privacy* are considered orthogonal to service quality and are usually separated from the service quality description with respect to the other security aspects. In this way, another partition of SQMMs is considered in this section which maps to languages, such as Trust-Serv [Skogsrud et al. 2004], PeerTrust [Nejdl et al. 2004], WS-Trust [Nadalin et al. 2007], and P3P [Cranor et al. 2006], that describe a service quality part which is not described in the rest of the SQSLs. The SQMMs of this type are called *security-based SQMMs*.

4.2. Methodology and Analysis

In order to analyze all SQMMs and compare them on their ability to define quality, particular comparison criteria have been chosen that were either devised by the authors or collected from other research approaches [Frølund and Koistinen 1998; Tosic et al. 2002; Maximilien and Singh 2002; Cortés et al. 2005; De Paoli et al. 2008; Kritikos 2008; Kritikos and Plexousakis 2009]. These criteria mainly reflect the formality, expressiveness, complexity, and applicability of the examined SQMMs. The summary of our selected criteria (without their grouping) is shown in Table IV, while their complete presentation is provided later on.

Based on the SQMM categorization of the previous subsection, there are actually three SQMM partitions: *pure*, *SLA enabled*, and *security based*. The content of these partitions with respect to the corresponding SQMMs can be seen in Table V. The evaluation results of the examined SQMMs according to our comparison criteria are presented in Table VI. This table is separated into three clusters, each one presenting the evaluation results of one particular SQMM partition, where the respective approaches are ordered chronologically.

In the remainder of this section, each criterion or group of criteria is presented along with its evaluation results in separate subsections, where the analysis of the evaluation results takes place both globally for all SQMMs and locally for each partition. At the end, a global analysis of the SQMMs across all criteria is given.

Table V. SQMM Partitions and their Corresponding Approaches

Partition	Approach Name	Approach Reference
Pure	WS-QoS	[Tian et al. 2003]
	WSAF-QoS	[Maximilien and Singh 2004]
	DAML-QoS	[Zhou et al. 2004]
	QoSOnt	[Dobson et al. 2005]
	QRL	[Cortés et al. 2005]
	UML QoS	[The OMG Group 2005]
	WSMO-QoS	[Wang et al. 2006]
	OWL-Q	[Kritikos and Plexousakis 2006]
	onQoS-QL	[Giallonardo and Zimeo 2007]
	PCM	[De Paoli et al. 2008]
SLA-enabled	QML	[Frølund and Koistinen 1998]
	WSOL	[Tosic et al. 2003b]
	WSLA	[Keller and Ludwig 2003]
	SWAPS	[Oldham et al. 2006]
Security-based	Trust-Serv	[Skogsrud et al. 2004]
	PeerTrust	[Nejdl et al. 2004]
	P3P	[Cranor et al. 2006]
	WS-Trust	[Nadalin et al. 2007]

4.2.1. Formalism. This criterion has been chosen in order to distinguish SQMMs (i.e., the metamodels that define the abstract syntax of quality languages) depending on their representation formalism. Various formalisms have been used to express an SQMM, such as *informal* (such as DTDs or XML schemas), *UML*, and *ontologies*. Each formalism has its own advantages and disadvantages. For example, ontologies provide a formal, syntactic, and semantic description model of concepts, properties, and relationships between concepts. They are extensible, human understandable, and machine interpretable and enable reasoning support by using Semantic Web technologies. However, sometimes their expressive power is more than needed while also the tool support is not as efficient as in the other formalisms. Moreover, current ontology tools from the research community require expertise in knowledge representation. The evaluation of this criterion for each SQMM can have one of the following values: *informal*, *UML*, and *ontologies*.

The evaluation results are presented in the second column of Table VI. Most of the approaches use either ontologies or informal formalisms (mostly schema languages focusing on a language's concrete syntax), while only two use UML. This result is reasonable as ontologies are powerful modelling formalisms and very expressive, while informal formalisms are simple and very well supported. Concerning the local partitions, pure SQMMs are modelled mostly through ontologies, while an informal formalism is the best modelling choice in the other partitions. For pure and SLA-enabled SQMMs there is a recent trend to use ontologies for their modelling. This is because ontologies provide unambiguous semantics to quality terms and thus enable machines to automatically process and reason on ontology-specified QSDs to support service life-cycle activities like discovery and negotiation.

4.2.2. Richness in Defining Service Quality Models. In the presence of multiple services with overlapping or identical functionality, SRs need objective quality criteria to distinguish between these services. However, it is not practical to create a standard SQM that can be used for all services in every domain. This is because quality is a broad concept that encompasses many nonfunctional properties, such as privacy, reputation, and usability. Moreover, when evaluating service quality, domain-specific criteria must be considered.

Table VI. Comparison of SQMM Approaches on the 11 Evaluation Criteria

Research Approach	Formalism	Quality Model	Metric Model	Constraint Descr.	Com/xity	Descr. Sep., Refin., Granul.	Symmetric but Separate QoSs	Class of Service	Connection	Quality Matching Constructs	Framework Support
WS-QoS	Informal	Fair	Fair	Poor	Low	[n,y]	Symmetric	Many	WSDL	No	No
WSAF-QoS	Ontologies	Good	Poor	Poor	Low	[n,n,n]	Same Construct	One	OWL-S	No	Discovery
DAML-QoS	Ontologies	Fair	Fair	Rich	Low	[y,y]	Symmetric	Many	OWL-S	No	Discovery
QoSOnt	Ontologies	Fair	Fair	Good	Low	[y,y]	Symmetric	Many	OWL-S	No	Discovery
QRL	Informal	Poor	Fair	Rich	Low	[y,n,n]	Symmetric	One	No	No	Discovery
UML QoS	UML	Good	Fair	Rich	Medium	[y,y]	Same Construct	Many	No	No	No
WSMO-QoS	Ontologies	Fair	Good	Good	Low	[y,y,n]	Same Construct	Many	WSMO	No	Discovery
OWL-Q	Ontologies	Good	Rich	Rich	High	[y,y]	Symmetric	Many	OWL-S	Yes	Discovery
onQoS-QL	Ontologies	Good	Good	Rich	Medium	[y,y]	Symmetric	Many	OWL-S	Yes	Discovery
PCM	Ontologies	Fair	Good	Rich	Medium	[y,y,n]	Symmetric	Many	WSMO	Yes	Discovery
QML	Informal & UML	Fair	Good	Excellent	Low	[y,y]	Same Construct	One	No	No	No
WSOL	Informal	Poor	Good	Rich	Low	[y,y]	Same Construct	Many	WSDL	No	Discovery
WSLA	Informal	Poor	Good	Excellent	Medium	[y,y]	Same Construct	Many	WSDL	No	Negotiation
SWAPS	Ontologies	Fair	Good	Rich	Medium	[n,y]	Same Construct	Many	WSDL	Yes	Discovery
Trust-Serv	Informal	Fair	Fair	Good	Low	[y,y]	Same Construct	Many	WSDL	No	Discovery
PeerTrust	Ontologies	Good	Good	Rich	Low	[y,y]	Same Construct	Many	OWL-S	No	Discovery
P3P	Informal	Good	Fair	Good	Low	[y,y]	Same Construct	Many	WSDL	No	Discovery
WS-Trust	Informal	Fair	Fair	Poor	Low	[y,n,n]	Same Construct	One	WSDL	No	No

Therefore, a rich and extensible SQM must be enabled by the SQMM that includes both generic and domain-specific attributes and can be extended with the addition of new attributes. All the SQMMs considered advertise that they are extensible according to this aspect.

However, an SQMM richness depends also on other additional modelling capabilities that concern the offering of constructs that enable a quality attribute description in every possible detail/aspect. The following list summarizes all the modelling criteria that are needed to assess an SQMM's richness in defining SQMs:

- (1) enumeration of all possible quality attributes;
- (2) modelling the attribute's domain (e.g., phone service provisioning) (i.e., the entity and its relation to the "attribute" entity);
- (3) modelling of interattribute relationships/dependencies (either quantitative or qualitative or both);
- (4) modelling the attribute's compositionality (i.e., if it is composite or not and what are its child attributes);
- (5) modelling the different views which an attribute may concern, that is, the SPs, SRs, or both views;
- (6) distinguishing by using appropriate constructs between QoS and QoE attributes;
- (7) distinguishing by using appropriate constructs between domain-dependent and domain-independent attributes;
- (8) modelling the service layer an attribute refers to;
- (9) modelling the association/relationship between quality attributes and metrics.

It must be noted that for the last subcriterion we do not inspect the richness of the metric model, as this is the subject of the next criterion.

This criterion's evaluation for a specific SQMM depends on the SQMM's satisfaction of the nine previously analyzed subcriteria. In particular, the SQMM space is partitioned in a uniform way depending on the number of subcriteria that are satisfied by an SQMM. Therefore, if the considered SQMM satisfies only one or two subcriteria, it is considered *poor*. If it satisfies three to four criteria, it is considered *fair*. If it satisfies five to seven criteria, it is considered *good*. Finally, if the metamodel satisfies eight to nine criteria, it is considered *rich*.

The evaluation results are presented in the third column of Table VI. Most of the SQMMs can describe either a *fair* or *good* in richness SQM, while the rest can only describe a poor SQM. As there is no SQMM that can describe a rich SQM, there is no approach perfectly capturing this modelling aspect. Concerning the local partitions, the aforesaid general result also applies to pure SQMMs, while there is a trend revealing that the SQM richness is improved in recent pure SQMM approaches. Moreover, SLA-enabled SQMMs do not perform well in this modelling aspect because they place more importance on how a quality attribute can be measured rather than how it is modelled. On the other hand, security-based SQMMs perform moderately as the definition of quality attributes is one of their major concerns but not all attribute aspects need to be described.

4.2.3. Richness in Defining Quality Metric Models. An attribute is measured through the abstraction of a metric. While a metric model is encompassed in a quality model, the capability of an SQMM to express such models in a rich way was inspected separately for two main reasons. First, quality attributes and metrics are two different concepts. Second, both of these concepts require a quite rich description as the number of subcriteria that were used to compare the SQMMs' richness in describing both of these concepts can reveal. So, the overall richness of an SQMM depends on the richness of both its quality and metric models.

Thus, SQMMs should enable the creation of rich service quality metric models. The richness of the SQMM metric models was evaluated according to the following metric aspects.

- (1) *The metric dynamicity.* Metrics should be distinguished between dynamic and static ones. Dynamic metrics measure dynamic quality attributes (like *availability*) that change over time and are computed according to a schedule or trigger. Static metrics measure static quality attributes (like *security*) and have a specific value that does not change over time. This means that these attributes are not only controllable but also fixable, so SPs can guarantee their fixed values even if the provided service's context changes.
- (2) *The metric value type.* A metric value type specifies a domain of values. These values can be used in constraints involving this metric. The domain of values may be ordered. For example, a numeric domain comes with a built-in ordering that corresponds to the usual ordering on numbers. So, only the maximum and minimum value along with their numeric type (e.g., real or integer) have to be modelled for numeric domains. If the domain is not continuous, then it can be expressed as a union of continuous domains. In practice, numeric domains are used for most quality metrics. Set and enumeration domains do not have a built-in ordering, so a user ordering of the domain elements must be described apart from the explicit modelling of these elements. Depending on the quality metric semantics (e.g., if the amount of values the metric can take only matters), the natural partial order of sets defined by inclusion can be used. The order in which an enumeration's elements are defined may also be their sorting order but the user has to define if it is increasing or decreasing. Finally, for unordered domains only the domain values have to be explicitly defined.
- (3) *The metric unit.* The values a metric can take are measured in specific units, for example, *seconds* for a metric measuring *execution time*. Concerning the modelling of units, describing just the unit name is not enough because additional information regarding how to convert a value expressed in a specific unit to a value expressed in another unit has to be modelled. This information is crucial in case the SP and SR express a metric constraint using different units or when combining metric measurements of different units originating from different sources. Units should be separated into *basic* and *derived* units. Basic units should have a name and a short abbreviation. Derived units are produced from other units by multiplying the component unit with a specific (float) value (i.e., multiples of units) and possibly dividing it with another one. For example, the unit of *minutes* is produced by multiplying the unit of *seconds* with 1/60, while one unit for the *throughput* quality attribute is "bytes/second" produced by dividing the unit of "bytes" with the unit of "seconds". Thus, two relationships, stating which unit(s) are directly proportional and which are inverse proportional to it, and an additional multiplying factor should be modelled for a derived unit.
- (4) *The metric measurement directive or function.* Quality metrics should be classified into *resource* and *composite* metrics. *Resource metrics* are directly measured from the service's system instrumentation through *measurement directives*. For measurement directives, a URI, specifying how the value of a managed resource is retrieved, and the value type of the return value should be described. In addition, the *access model* (i.e., *push* or *pull*) must be specified to establish if the party responsible for the measurement will ask for the value or receive it when it is ready. Moreover, specific measurement directives may require a possible extra attribute (timeout) specification concerning the time duration that the measurement party will wait for obtaining the measurement value. *Composite metrics* are computed by

applying statistical or other mathematical functions to other metrics. Therefore, a description of both the function and the metrics used to compute the composite metric must exist. Moreover, a function model should be provided in order to enable users to select the appropriate function for each composite metric.

- (5) *The metric schedule.* The SQMM should enable the definition of at least one of the following types of time windows for the periodic or instantaneous calculation of new values for a quality metric:
 - (a) calendar time window like week, month, and/or year;
 - (b) sliding windows like the last ten days;
 - (c) expanding window or running total, for example, from this year's start until now.
- (6) *The metric weight relative to its implicit domain and user preferences.* This weight can be used to calculate the rank of a service quality offer and indicates the impact that this metric has to the overall quality offered by a service.
- (7) *The characteristic of the function from metric values to overall quality values.* An SQMM should explicitly specify the exact monotonicity/value direction of monotonic metrics (e.g., negative for an *execution time* metric) and mapping functions for nonmonotonic metrics. This information modelling is sufficient in most of the quality-based service discovery scenarios.

The nonmonotonic metrics need user-defined mapping functions to express the user preferences regarding the values that these metrics can take. For example, for the nonmonotonic metric enumerating the encryption algorithms supported by a service, there may be a user function that maps the value of "AES-192" to service level 3, the value of "AES-256" to service level 2, and the value of "AES-128" to service level 1. In this case, the highest security value gets a lower quality value than the second highest one indicating that the user may be satisfied with the second highest value and does not want to pay more in order to have a more secured service.

The preceding situation is tailored for QBSM scenarios where each quality metric is considered independently of the other. In case that there are dependencies between quality metrics and attributes, functions (or n-ary constraints) should be used to capture them in conjunction with the aforementioned mapping functions. The simple additive weighting technique [Hwang and Yoon 1981] is commonly used in service selection and requires that the values of each attribute or metric are normalized according to a specific evaluation function. In this case, the aforesaid mapping functions of nonmonotonic metrics can be used provided that they map the metric values to the same range (usually the [0.0,1.0] range). For monotonic metrics, particular evaluation functions are used in most research approaches that do not have to be captured in an SQMM.

- (8) *Aggregation of the values of a composite service's metric.* There must be a description (mathematical or otherwise formal) of how, for a complex service, the value of a quality metric can be calculated from the values of the corresponding quality metrics of the individual services it contains. For example, if T_c is the execution time of a complex service C , a sequence of two services A and B , then T_c can be computed as the sum of the execution times of the two individual services. This description is essential for the automated estimation of the quality metric values for a composite service. So this description is needed for automating the quality analysis process, a prerequisite for a successful quality-based service discovery.

This criterion's evaluation for a specific SQMM depends on the SQMM's satisfaction of the eight previously analyzed subcriteria, where, as in the case of the previous criterion, the SQMM space is partitioned in a uniform way. If an SQMM does not

comply with all the modelling requirements for a specific subcriterion, then this SQMM does not satisfy the subcriterion. So if the considered SQMM satisfies only one or two subcriteria, then it is considered *poor*. If it satisfies three to four criteria, it is considered *fair*. If it satisfies five to six criteria, it is considered *good*. Otherwise, if it satisfies seven to eight criteria, it is considered *rich*.

The evaluation results are presented in the fourth column of Table VI. Most SQMMs encompass either a good or a fair metric model, while only one SQMM encompasses a poor model. In addition, only one approach (OWL-Q) captures a rich metric model, while it does not offer a rich quality model. Thus, a metric model is better captured in the current state-of-the-art SQMMs than a quality model. Considering the local results in each partition, SLA-enabled SQMMs capture a good metric model with respect to its richness. This observation was actually justified in the previous criterion's evaluation, where it was stated that SLA-enabled approaches pay more attention to metric than to attribute modelling. The majority of pure and security-based SQMMs present a fair metric model. Finally, there is a trend for pure SQMMs to improve their metric model as the SQMMs proposed after 2004 do not encompass a fair or poor metric model. By combining this result with the respective result of the previous criterion, it can be deduced that the most recent pure SQMMs have increased their expressiveness as they cater for a better service quality and metric model.

4.2.4. Expressiveness in Constraint Description. QSDs comprise quality constraints, which consist of a name, an operator, and a value [Frølund and Koistinen 1998]. The name is typically a quality metric's name, or the name of a metric *aspect* or *function*. The permissible operators and values depend on the quality metric's value type. Only the values that a metric can take should be used in constraints for that metric. The domain may be ordered. The domain ordering determines which operators can be used in constraints for that domain. For example, only the equal "=" and unequal operators "≠" can be used in unordered domains and not inequality operators ("<", ">", "≤", "≥").

Aspects [Frølund and Koistinen 1998] are statistical characterizations of quality constraints, such as *percentile*, *mean*, *variance*, and *frequency*. They are used for the characterization of measured values over some time period. For example, the percentile aspect could be used to define an upper or lower value for a percentage of the measurements or occurrences that have been observed. Aspects can be very useful when the measurements or occurrences of a quality metric present some special characteristics and a new complex metric should not be produced from the basic quality metric for each of these characteristics.

Quality constraints are usually connected by logical operators into expressions. A service quality specification should contain one complete constraint expression or just one constraint. Moreover, quality constraints should be joined into *Constraint Groups* (CG) or *Constraint Group Templates* (parameterized CGs) in order to be reused by many service quality specifications [Tosic et al. 2003a]. *Constraint Groups* contain a set of *concrete* quality constraints, while *Constraint Group Templates* contain *abstract* quality constraints (i.e., the second constraint operand is not specified). Other reusability constructs can also be created even for expressions.

To evaluate this criterion for a specific SQMM, the following cases are considered, which depend on whether the SQMM satisfies any of the aforesaid requirements, where these requirements have been separated in different levels of importance/significance. In this way, if the SQMM does not satisfy any of the previous requirements, then its expressiveness is *poor*. If the SQMM allows comparison operators in constraints but does not check their compatibility with the metrics used, then its level is increased and the SQMM is considered *fair*. If the SQMM allows comparison operators and checks their compatibility with constraints, then it is promoted at the *good* level. If the SQMM

is *good* and allows not only to express constraint conjunctions but also other constraint expressions and formations, it is placed at the *rich* level. Finally, if the SQMM is *rich* and allows the specification of aspects, then it is placed at the highest possible level, namely the *excellent* one.

This criterion's evaluation results are presented in the fifth column of Table VI. Most of SQMMs use a rich constraint model. As there are no fair constraint models, SQMMs either perform above or below the average in this aspect. Moreover, in the SLA-enabled partition, there is a balance between those SQMMs that capture excellent constraint models and those that capture rich. This is because the rich constraint representation, especially for service quality levels, is one of the cornerstone features of SLA languages, which is equally important for all SLA contracting parties. The security-based SQMMs mostly use a good constraint model because they do not require advanced constraint modelling features. In pure SQMMs rich constraint models are the majority. Moreover, pure SQMM modelers have quickly understood this feature's importance as their SQMMs moved from a poor to a better constraint model. By combining this result with the respective results of the two previous criteria, it can be deduced that pure SQMMs have increased their overall expressiveness over time.

4.2.5. Complexity. On one hand, a metamodel's *formalism* characterizes the explicitness in which the semantics of the metamodel's terms are expressed, while its *expressiveness* or *richness* concerns how well the domain of discourse (i.e., service quality) is modelled. On the other hand, a metamodel's complexity mainly concerns its size and structure and significantly impacts user understandability, the modelling effort involved, and the size and amount of information included in the produced descriptions. Obviously, the better the domain is modelled and more details are captured, the higher is the metamodel complexity. So, usually there is a trade-off between complexity and richness that is mainly regulated by the metamodel's quality. The latter depends on how accurately and extensively the metamodel expresses its domain, the existence of formal and semantic inconsistencies in it, and the relevance and appropriateness of the modelled information. Thus, rich and qualitative metamodels must be developed such that their complexity is not big.

It is very difficult to assess a metamodel's quality and very few approaches have been proposed focusing on some quality aspect [Jiang et al. 2004; Welty et al. 2003]. On the contrary, many metrics have been proposed for measuring metamodel or ontology complexity [Mens and Lanza 2002; Yi et al. 2004; Yoa et al. 2005; Ma et al. 2010]. Thus, the SQMM quality can be speculated by assessing the SQMM complexity through one of these metrics in combination with the SQMM richness.

The complexity metric to measure a metamodel's number of concepts/elements was chosen for the following reasons: (a) many SQMMs are not publicly available, so they cannot be evaluated with sophisticated complexity metrics; (b) many complexity metrics are not applicable to informal SQMMs; (c) the structure particularities of the informal SQMMs require using a simple and fair metric that does not depend on the SQMM structure and is easy to compute. However, such a simplistic metric prevents performing reasonable speculations about the SQMM quality, as the threshold on the number of sufficient concepts for modelling a domain depends on the subjective view of domain modelers.

Particular thresholds on the number of concepts involved in SQMMs were used to assess their complexity, which were set in a uniform way according to the lowest and highest number of concepts across all SQMMs. So, if an SQMM has less than 25 concepts, it is evaluated to have *low* complexity. If it has 26 to 50 concepts, then it has *medium* complexity. Finally, if it has more than 50 concepts, then it has *big* complexity.

This criterion's evaluation results are presented in the sixth column of Table VI. Most SQMMs have low complexity, while only five SQMMs have medium complexity. Only one SQMM (OWL-Q) has high complexity, which is the most expressive SQMM according to the evaluation of the previous three criteria. Concerning the SLA-enabled partition, there is a balance between the approaches with low and medium complexity. All security-based SQMMs exhibit low complexity. By inspecting the results, there is a trend in proposals for pure and SLA-enabled SQMMs of higher complexity. This means that modelers are increasing the expressiveness of their SQMMs and, in result, the complexity of their SQMMs. Indeed, based on the analysis of the previous criterion, pure SQMMs have increased their expressiveness over the years. Section 5 will show that the SLA-enabled SQMMs have increased their expressiveness but in SLA-based aspects and not in quality-related ones. Security-based SQMMs have only low complexity as they are deliberately designed in this way, that is, to produce short descriptions that are easily exchanged, processed, and matched by various entities over the Web.

4.2.6. Service Description Separation, Refinement, and Granularity. This subsection presents and analyzes a group of three related subcriteria. As these three criteria represent features that an SQMM either possesses or not, the SQMMs are evaluated with either y (possess the feature) or n (do not possess it) for each criterion. This results in an overall evaluation result for an SQMM of $[t_1, t_2, t_3]$, where $t_i = y$ or n and the evaluation values order is the one with which the criteria are analyzed next.

4.2.6.1. Syntactical separation of quality-based and functional parts of service description. Service quality specifications should be syntactically separated from other parts of service specifications, such as interface definitions. This separation allows to specify different quality properties for different implementations of the same interface. Moreover, while functional constraints rarely change during runtime, service quality constraints do change. So, the separation of service quality offers from WSDL descriptions permits these offers to be deactivated, reactivated, created, or deleted dynamically without any modification of the underlying WSDL file. Finally, a service quality offer could be referenced from multiple WSDL files and thus be reused for different services.

4.2.6.2. Support for the refinement of quality-based service descriptions and their constructs. Apart from syntactical separation, another form of reusability is equally important. Service quality specifications should be also refined. This means that a new service quality offer can be created by referencing an older one and adding constraints like the refinement of an older quality restriction or creation of a new one. In addition, templates of service quality offerings should be created and appropriately extended for every domain.

4.2.6.3. Support for fine-grained quality-based service description. It should be possible to specify quality properties/metrics at a fine-grained level. As an example, performance characteristics are commonly specified for individual operations. An SQMM must allow quality specifications for interfaces, operations, attributes, operation parameters, and operation results.

This criterion's evaluation results are presented in the seventh column of Table VI. Most of SQMMs enable all three features. So, most SQMM modelers have understood the importance of these three features. The same result holds for all the partitions. The mostly supported features among all SQMM approaches are syntactical separation and specification refinement, followed by fine-grained specification. The same order holds for the pure and security-based SQMMs. Finally, for SLA-enabled SQMMs the order is different. Fine-grained specification is the most important feature, while the other two are of equal importance.

4.2.7. Support for Symmetric but Separate Quality-Based Service Description for Providers and Requesters. Both the quality properties that clients require and the quality properties that services provide must be specified separately so that a client-server relationship has two service quality specifications: a specification that captures the client's requirements (i.e., service quality requirement) and a specification that captures service provisioning (i.e., service quality offer). This separation allows us to specify the quality properties that a component provides or requires, without specifying the interconnection of components in order to enable the component reuse in different contexts.

Service quality requirements and offers should be specified in the same expressive way, that is, symmetrically. Assuming that S is a multidimensional space whose dimensions are given by domains of quality parameters, then both a service quality offer and requirement should be expressed as a subspace in S . Traditionally, service quality offers have been described as points in S , that is, asymmetrically. However, this semantics makes it difficult to specify offers whose quality can vary. Moreover, the probability that an offer is matched with a requirement is quite low. On the other hand, the probability of matching is higher when both types of QSDs are expressed as subspaces, while also more advanced protocols are enabled in service negotiation. In addition, symmetric approaches achieve a better expressiveness to specify QoS, since there is usually no restriction on the number of involved parameters or operator types, so that nonlinear or more complex expressions are allowed.

This subcriterion inspects if both SPs and SRs can provide service quality specifications, if these specifications are defined separately with different constructs, and if these specifications are allowed to have the same expressiveness. This criterion's evaluation for a specific SQMM has the following values: (a) *SP*, (b) *same construct*, (c) *asymmetric* (so also different constructs are used), and (d) *symmetric* (where different constructs are used). The first and last values are the worst and best, respectively, while there is no definite order between the second and third value as each violates a different requirement.

The evaluation results are presented in the eighth column of Table VI. Most SQMMs allow both SPs and SRs to specify service quality specifications with the same construct, followed by those SQMMs that allow both SPs and SRs to specify service quality specifications with different constructs but symmetrically. Thus, researchers have realized that both SPs and SRs should be allowed to specify their QSDs with the same expressiveness. Concerning the SQMM partitions, all SLA-enabled and security-based SQMMs use the same construct for expressing the two QSD types. This is a rational choice for SLA-enabled SQMMs as the produced SLAs express both the views of the SP and SR. However, for security-based SQMMs, this is not a rational choice for privacy, as there must be a clear distinction between who is requesting privacy requirements and who is offering to satisfy these requirements, but it is rational for trust, as the specification of requirements and offerings is performed in a bilateral way. The majority of pure SQMMs express the two different QSD types in a separate and symmetric way. Thus, the pure SQMM modelers have realized the significance of this design choice.

4.2.8. Support for Classes of Service Description. *Class of Service* means the discrete variation of the complete service and quality provided by one service. In our opinion a *Class of Service* has the same meaning as the *alternatives* have in WS-agreement. In other words, a *Class of Service* is actually the set of (functional and) nonfunctional guarantees that are to be provided by a service in terms of an SLA. So, this criterion assesses if a SQMM can capture only one class of service specification or many. In this way, this criterion's evaluation for a specific metamodel is: *one* (class of service), and *many*.

The evaluation results are presented in the ninth column of Table VI. Most SQMMs can represent many classes of service. This result also applies to each partition

separately. So SQMM modelers have understood the advantage of allowing many classes of service specification for one service. For pure and SLA-enabled SQMMs this understanding is more anticipated in the most recent approaches.

4.2.9. Connection to Functional Service Specification Languages. This criterion inspects if the SQMM is connected to or references a Service Functional Specification Language (SFSL) like WSDL, WSMO, or OWL-S. On one hand, if the answer is positive, this means that there would be no effort required to extend this SQMM (and probably all of its QSDs) in order to be used in a registry containing functional SDs obeying the connected SFSL. On the other hand, if the answer is negative, then the SQMM can be extended in order to be connected to any SFSL and not only to a specific one. In this way, it can be practically used with any registry.

The evaluation results are presented in the tenth column of Table VI. WSDL is the most referenced SFSL, followed by OWL-S (ontology-based SFSL), while another ontology-based language, namely WSMO, is referenced by only two SQMMs. Three SQMMs do not reference any SFSL, being in this way SFSL independent. Taking specific SQMM partitions into consideration, OWL-S (or ontology-based SFSLs more generally) is the most referenced SFSL in pure SQMMs for two main reasons: (1) ontological approaches are greater in population with respect to the rest of the approaches in this partition, and (2) the use of semantics has been proven to increase the accuracy in FSD, so pure SQMM modelers prefer an ontology-based SFSL for this reason. On the other hand, WSDL is the most referenced SFSL in the rest of the SQMM partitions. This is expected because SLA-enabled and security-based SQMMs do not consider the FSD scenario when they are designed, as they regard it as an orthogonal issue. So they prefer to stay on the most popular and used SFSL (i.e., WSDL) in order to increase their adoption. Moreover, there is a trend that the most recent SLA-enabled SQMMs are connected to WSDL.

4.2.10. Support for Quality Matching. As users may have different conceptualizations of the same domain, it is possible that in different QSDs of the same SQMM the same concepts are expressed differently or different concepts are expressed in the same way. Concerning the QoS domain, this can be true for QoS metrics and QoS attributes but not for measurement units that are more or less standardized. This argument is also strengthened by the fact that the same QoS metric or attribute can be considered either composite or atomic in different SQMs depending on the level-of-detail required or the measurement types supported. For instance, a service's *availability* is measured from high-level readings in one system's instrumentation, while it is measured from low-level readings (e.g., from service's *uptime*) in another system's instrumentation. In the former case, the service's availability metric is resource, while in the latter case, the same metric is composite.

As the basic QSD part is the one where QoS capabilities or requirements are expressed as a set of constraints on specific QoS metrics or attributes, these sets of QoS concepts must be matched in order to increase the accuracy of the QBSM process. Thus, semantic QoS metric/attribute matching rules must be developed and used internally or externally to an SQMM in order to enrich and align the produced QSDs. Depending on the SQMM's formalism, rules may be part of the SQMM's specification or may be developed externally in a possibly different format. In the latter case, it could be argued that these external rules are not actually part of the SQMM's specification and, thus, they should not be a criterion for comparing SQMMs. However, this information, either being internal or external, should be additionally modelled as it can be used to support service discovery and increase an SQMM's adoption and significance. Thus, this criterion's evaluation for each SQMM will be: *yes*, if the corresponding SQMM contains or is accompanied with quality matching rules and *no* otherwise.

The evaluation results are presented in the eleventh column of Table VI. Most of the SQMM approaches have not considered this modelling aspect at all. This also holds locally in each partition. Moreover, not all security-based SQMMs use or model these matching rules as they are not required because the quality attributes in this area are more or less standardized, while quality metrics are used less often. Finally, the most recent pure and SLA-enabled approaches have understood the need of modelling this feature and have produced the required matching rules.

4.2.11. Framework Support. This criterion inspects if any service discovery and negotiation framework has adopted the SQMM under inspection. If the answer is positive, then this means that the SQMM has been used in practice in one of the service life-cycle activities (apart from the first one). So this criterion evaluates the adoption and usage of every SQMM. In addition, it evaluates if SQMMs have been used in service negotiation apart from service discovery. The evaluation of a specific SQMM for this criterion is: *no* (support), *only discovery*, *only negotiation*, and *both* (discovery and negotiation). As all SLA-enabled SQMMs already have an underlying SLA enforcement framework implementation, we chose not to explicitly represent this fact in this evaluation.

The evaluation results are presented in the last column of Table VI. The majority of the SQMMs have been used in service discovery implementations. This result holds not only globally but also in each partition. In addition, in conjunction with the results of the previous criterion, all the SQMM approaches that model matching rules have been used in service discovery implementations, thus realizing the need for increasing the accuracy of the service discovery activity. However, none of the implementing discovery frameworks can also perform service negotiation. Moreover, there is no service negotiation framework supporting any of the SQMMs. It must be noted that there are negotiation frameworks supporting WS-agreement but not SWAPS (which is a WS-agreement extension).

Most of the SLA-enabled SQMMs apart from WSLA are not used in any implemented negotiation framework. This actually reduces the dynamics and further spread of these SQMMs. In addition, by considering the service life-cycle analyzed in Section 2, it should be stated that only the SLA-enabled SQMMs have been used in service negotiation, as the QSDs of this SQMM type that are used in this activity are SLAs or SLA templates. Thus, the service negotiation activity is the stopping usage point of the QSDs and the corresponding starting point of SLAs.

4.3. Overall Analysis

Based on the previous analysis, there is not any complete approach that has taken the best value in all criteria. From all the SQMMs, OWL-Q can be distinguished as the best among the good approaches which are found mainly in the pure SQMM partition. However, it must be extended accordingly so as to enable the specification of rich SQMMs and even more expressive constraint models.

Considering the remaining SQMM partitions, SWAPS is the best among all SLA-enabled SQMMs but it must significantly improve its attribute and metric model. This is also true for PeerTrust, the best among all security-based SQMMs. Thus, the SLA-enabled and security-based SQMMs need further improvement and extension in their attribute and metric models and far more changes with respect to those needed for the pure SQMM approaches.

An ideal SQMM that satisfies most of the comparison criteria is visualized as an UML diagram in Figure 4. The parts that are highlighted with the grey color correspond to those conceptual elements that are common among the majority of all SQMMs. These elements have been identified by following a similar approach as the one of Section 3.3 which was used to identify the most common quality attributes among

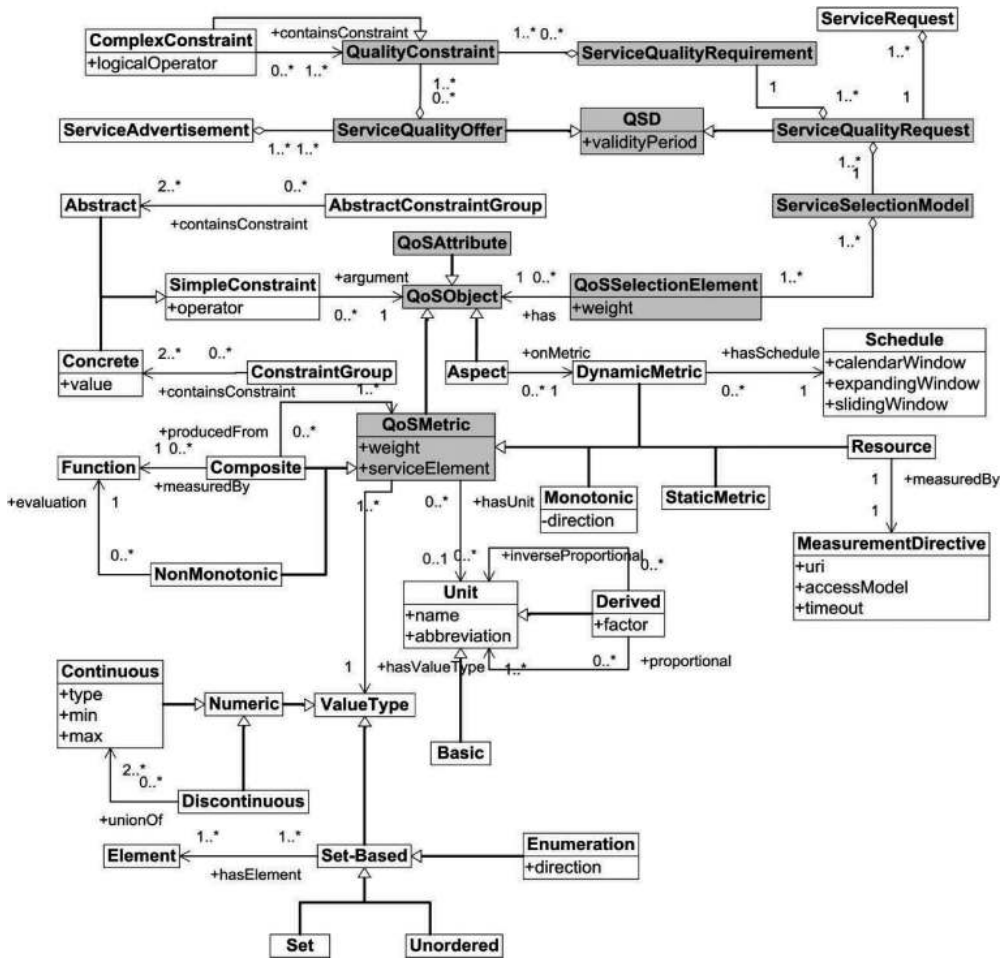


Fig. 4. An ideal SQMM.

SQMs. The *QoS Attribute* concept is not further analyzed, as an ideal metamodel related to the appropriate and rich quality attribute description has been already presented in Figure 3 of Section 3. This latter metamodel is part of the ideal SQMM presented.

4.3.1. Ontology-Based SQMMs. As indicated previously, there is now a trend in adopting ontology-based SQMMs due to the advantages that their usage brings up. To this end, let us now focus on this type of SQMMs by unveiling some interesting facts about their content as well as analyzing their quality.

One interesting fact that can only be inferred by inspecting all ontology-based SQMMs is that some of them define concrete quality attributes and metrics at concept level (e.g., DAML-QoS), while the others (e.g., OWL-Q) define such terms and their relationships at instance level. The rationale for the definition at the instance level relates to the fact that the latter SQMMs use specific rules that can be used to infer the equivalence between quality terms, to align QSDs, as well as to match or select QSDs [Kritikos 2008], while the definition of these terms at the concept level enables equivalence and subsumption relationships between terms to be unveiled as well as the matching of QSDs [Zhou et al. 2004].

Table VII. The Cohesion and Complexity Assessment Results of the Ontology-Based SQMMs

Approach	NOC	NOL	NOP	NOF	AC-F	MAX-DIT	ADIT-LN
WSAF-QoS	19	13	8	22	1,16	3	2,42
DAML-QoS	13	10	8	10	0,77	2	2
QoSOnt	22	11	21	10	0,45	2	2
WSMO-QoS	4	3	10	1	0,25	2	2
OWL-Q	107	70	82	169	1,58	5	2,98
onQoS-QL	44	26	39	41	0,93	3	2,46
PCM	26	14	12	32	1,23	3	2,86
SWAPS	26	6	28	6	0,23	2	2
PeerTrust	13	5	15	5	0,38	2	2
Ideal SQMM	49	24	53	53	1,08	4	2,91

Evaluating the quality of an ontology has received great attention for many years and many related approaches have been proposed. The most prominent from these approaches [Yoa et al. 2005; Ma et al. 2010] focus on particular aspects of ontology quality, such as complexity and cohesion and have proposed formal sets of metrics which can be used to assess an ontology based on these aspects. An ontology is cohesive if its entities are strongly related to each other. In this sense, the cohesion of an ontology resembles the cohesion of software and the respective approaches aim at evaluating the separation of responsibilities and the independence of ontology components. On the other hand, the complexity of an ontology is potentially a measure of its completeness, as the more complex is the ontology, more and more appropriate domain concept and properties are captured and thus the ontology becomes more complete. However, if an ontology becomes very complex without being also cohesive, then this is a sign that there is a loss of focus, as, for example, more irrelevant concepts to a domain are modelled.

As many of the proposed ontology-based SQMMs are not publicly available, we just have a description of the modelled concepts and their attributes and relationships. This means that: (a) there is no information about any other type of axioms and (b) we cannot make any assumption about their consistency. Based on this setting, the most appropriate complexity and cohesion ontology metrics were drawn from the approach of Yoa et al. [2005] and are the following.

- NOC: Number of Concepts
- NOL: Number of Leaf concepts
- NOP: Number of Properties
- NOF: The sum of the fanout of all concept nodes
- AF-C: The average concept fanout
- MAX-DIT: The maximum path size from a leaf to a root concept
- ADIT-LN: The average path size from a leaf to a root concept

Before presenting the results of the comparison of the ontology-based SQMMs based on the preceding metrics, it must be indicated that to ensure the fairness of the comparison, the following restrictions were imposed: (a) only the upper ontology, which actually maps to the respective SQMM, of an approach was considered and not the mid- and low-level ontologies which map to concrete domain-independent and domain-dependent quality terms, respectively, and (b) unnamed classes were not considered. Moreover, to make the comparison even more interesting, the ideal SQMM was also considered. The results of the assessment of the complexity and cohesion of the ontology-based SQMMs are shown in Table VII.

As can be seen from the results presented in Table VII, the ontology-based SQMMs that exhibit good levels of complexity and coherence are only the pure-based ones and

especially the WSAF-QoS, OWL-Q, onQoS-QL, PCM, and the ideal SQMM. However, WSAF-QoS and PCM are not complex enough to capture all the appropriate aspects of service quality description. This was also indicated by the analysis of the second, third, and fourth comparison criteria in Section 4.2. The ideal and the onQoS-QL seem to have similar complexity (see first three criteria which are mostly related to an ontology's complexity) but the ideal SQMM seems to be slightly more complex and even more cohesive (see the last three metrics which are mostly related to an ontology's cohesiveness). However, the best ontology-based SQMM seems to be OWL-Q which is the most complex and cohesive ontology-based model. By relating these results with the comparison results of Section 4.2, it can be certainly concluded that OWL-Q is more complex due to its increased expressiveness with respect to the other SQMMs and that this complexity did not lead to a reduction of this ontology model's quality.

By comparing the complexity and cohesion of OWL-Q and the ideal SQMM it can be deduced that OWL-Q is more complex and slightly more cohesive. By also bearing in mind that this ideal SQMM could be further improved in some aspects, this means that the complexity of OWL-Q does not seem overwhelming. However, by also recording the fact that OWL-Q needs some improvement in some modelling aspects, this means that additional complexity will be attained, thus, a better balance between expressiveness and complexity should be accomplished in this SQMM by adding the appropriate concepts and relationships identified in Section 4.2 and possibly removing other terms which may be not very essential in service quality description. Another possible improvement concerns defining which are the obligatory and optional concepts in order to help modelers in defining their QSDs with less effort and more speed as well as supplying well-documented guides on how to produce complete QSDs. The latter improvement concerns not only OWL-Q but any other SQMM and if realized could lead to an increase in the SQMM's adoption and exploitation.

By closely inspecting the last criterion results, pure and security-based SQMMs are used until the service discovery activity and they seem not to be exploited further in the service life-cycle. This can be justified by their inability to model SLA-specific constructs which are considered more useful for supporting the rest of the service life-cycle activities. This should be the reason why SLA languages have been proposed, that is, to fill the gap and play the significant role of mostly supporting those service activities that are beyond service discovery. The next section explains this role and provides an analysis of the capabilities of the SLA languages with respect to their support to the service life-cycle activities.

5. SERVICE-LEVEL AGREEMENTS

5.1. Background

5.1.1. Contracts and SLAs. As the global economy is changing at a fast pace and the competition is rising due to technology advancements, organizations have to enter in dynamic business relationships with other organizations, whose result would be one or more cross-organizational processes. The basis for the establishment of such dynamic relationships are electronic *contracts*, which are legally binding digital agreements that safeguard the concerns of each participating organization [Hoffner et al. 2001], as they assist in speeding up and automating the various activities that support these relationships, which include the contractual relationships establishment and the enactment infrastructure setup. Many types of contracts exist based on the type of value that is exchanged between two or more organizations [Grefen and Angelov 2002]. However, this article considers only *service contracts*, as the focus is on services. A service contract includes information, such as [Hoffner et al. 2001; Oren et al. 2005]: (a) the *parties* involved in the agreement, (b) the *service*, including interface description and

expected interactions, (c) a description of the *norms* (e.g., *obligations*, *prohibitions*, *permissions*, etc.) imposed on each party concerning service provision and consumption, (d) timing and conditions of contract termination. Apart from technical information, some service contracts may also contain legal procedures in case of breach of contract and arbitration.

Most of the service contract languages have focused on describing the service functionality and on automated contract execution monitoring [Oren et al. 2005], that is, determining the state a contract is in, and which contract rules are in effect given this state. Thus, they focused more on how the service behaves while they more or less neglected the QoS aspects of service behavior. However, as QoS plays an important role in the whole service life-cycle and is equally important to functionality, the focus is now on more specialized service contracts, called SLAs. SLAs describe how well a service performs its functions [Tosic and Pagurek 2005]. They contain quality guarantees that have to be respected during service execution and other important terms indicating the actions to be taken when these guarantees are violated. Thus, they increase the trust and consolidate the overall relationship between an SP and an SR, as the service will either meet the stated requirements or there will be consequences that tend to compensate the client for the harm it suffers due to these requirements being missed [Skene 2007]. The last point is very important as it demonstrates how both service contracts and SLAs are the basis for the establishment and support of business relationships.

5.1.2. SLA Demystification. According to Paschke and Schnappinger-Gerull [2006], SLA documents contain *technical* (e.g., metrics, actions), *organizational* (monitoring and reporting), and *legal* (legal responsibilities, invoicing and payment modes) components. It is difficult to automate and enforce the legal components of SLA documents, so these are either omitted or neglected. The most common SLA components are the following [Paschke and Schnappinger-Gerull 2006]: *involved parties*, *contract validity period*, *service definitions*, *SLOs*, and *action guarantees*. *Involved parties* are roles referenced inside a contract. They can be distinguished between [Keller and Ludwig 2003] *signatory parties*, which are usually played by the SP and SR that sign the contract, and *supporting parties*, which have the role to support the SLA monitoring and assessment. These two role types are not mutually exclusive, as, for instance, an SP can provide measurements for the provided service's execution time. The *contract validity period* specifies for how long the SLA will be valid and enforceable. This field is important for *continuous* services as it determines the period that the SPs should provide their services to the requesters according to the directives of the agreed SLAs. The *service definitions* section specifies the service characteristics (i.e., functionality), components (i.e., operations, input, output, internal and external services for a composite service), and observable parameters (i.e., QoS metrics for the service and its components). *SLOs* are QoS guarantees that must be met by a specific obliged party (e.g., SP) and have validity periods [Keller and Ludwig 2003], while they can also have qualifying conditions on external factors such as time of the day (i.e., when the SLO should be evaluated) as well as the conditions that a client must meet (e.g., a client's request rate is above a threshold). Finally, *action guarantees* [Keller and Ludwig 2003] express a commitment that a particular activity is performed by an obliged party if a given precondition is met (e.g., a violation occurs). The committing activities include compensation, reward, recovery, and management actions.

Besides the common SLA components, two additional SLA components should be described [Muller 1999]: *limitations* and *renegotiation*. The former describes the limits of IT support during peak period demand conditions, resource contention by other applications, and general overall application workload intensities. These limitations should be agreed by all parties in order to prevent finger pointing and user dissatisfaction. The

latter describes how and under what circumstances the SLA must be changed through renegotiation to reflect changes in the user, service, or environmental context.

Before SLAs are established, they are in a form which is called *SLA template*.¹ These SLA templates are used to describe, matchmake, and negotiate the SLs to be offered by a service of an SP to an SR. Thus, they are produced by both SPs and SRs. SLA templates can be complete or incomplete SLAs. Complete SLA templates are commonly agreed among all participants in a restricted domain or are used as bilateral agreements between two organizations or as SLA offerings advertised by an SP to specific customer classes. Thus, they are offered in a “take it or leave it” basis [Hoffner et al. 2001]. Incomplete SLA templates can be seen as a skeleton with fields which must be completed according to the directives of the desired relationship between two organizations [Hoffner et al. 2001]. So, they are generic forms or templates that can be tailored to the specific circumstances of an SLA instance. According to the *granularity of choice* and *specialization* [Hoffner et al. 2001], they may: (a) be *monolithic* where values are inserted in predefined positions; (b) have certain sections which can be included or removed; or (c) be clause based [Fosbrook and Laing 1996].

There can be different SLA types according to their composability, intended purpose [Paschke and Schnappinger-Gerull 2006], and the service usage based on the reference functional architecture model [Lamanna et al. 2003]. Concerning the composability aspect, SLAs are distinguished between *single* and *composite*. *Single SLAs* specify the SLs of one service, independently of the service type (i.e., single or composite), and are agreed between two parties, the SP and SR. On the other hand, *composite SLAs* specify the SLs of both the service, which should be provided to an SR by the SP, and some of its components or supporting services that are provided by third-party SPs to the main SP. Thus, a *composite SLA* may consist or depend on other SLAs. This does not exclude the possibility that a *single SLA* may depend on one or more other SLAs. However, this dependency information is not included in the *single SLA*’s description.

5.2. Methodology and Analysis

As an SLA is a document, it has a life-cycle that starts with its creation and ends with its disposal or archiving. This life-cycle includes all the appropriate activities needed for the SLA management and is tightly coupled with the service life-cycle introduced in Section 2. The same holds for the contract life-cycle. This tight coupling is justified as follows. SLAs, and service contracts in general, exist as long as the service they describe exists because this service is the reason for the establishment and very existence of the business relationship between the SP and SR. Indeed, all service contract and SLA management systems actually support, directly or indirectly, the management of the service offered by the SP to the SR. Thus, through the support of the SLA/contract life-cycle, the service life-cycle is supported and especially those activities that correspond to service provisioning.

Service contract life-cycles (e.g., in Hoffner et al. [2001]), are more coarse-grained and general with respect to the SLA life-cycle as they consider both the service provisioning functional and QoS aspects. However, as the focus is on service quality and its description, only the QoS aspects of service provisioning are considered. Therefore, in this survey, the analysis of service contract and SLA languages is based on a set of comparison criteria which are grouped along the SLA life-cycle activities. These criteria are used to compare these languages along the lines of the information they can describe which is necessary for supporting the SLA life-cycle activities.

Various SLA life-cycles have been proposed in the literature, which differ on the activities they involve, the activities granularity level, and their terminology. However,

¹The term *Contract Template* is used for service contracts.

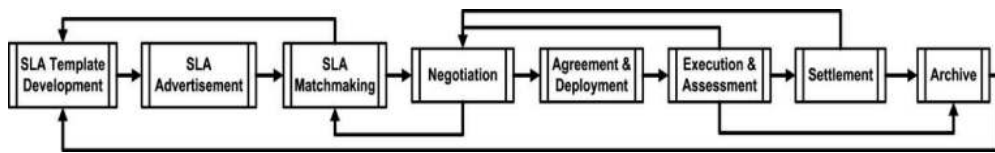


Fig. 5. The SLA life-cycle.

none of them is suitable in this article context as they are either coarse-grained or tend to neglect some activities. To this end, an extended SLA life-cycle has been devised based on the research works of Keller and Ludwig [2003] and Parkin et al. [2008], which is depicted in Figure 5.

In the following, using this life-cycle as a basis, the life-cycle activities involved in the SLA management are analyzed along with their relation to those of the service life-cycle.

- Template Development.* An SLA template is developed by the SP or the SR according to the service quality capabilities and requirements, respectively. This activity is executed after the service is implemented and tested but can happen before, during, or after the service deployment.
- Advertisement.* After the SLA template is developed by the SP, it is advertised in intra- or interorganizational repositories in order to be discovered by potential SRs. This activity happens after the service is deployed. It is a joint activity of the service and SLA life-cycles.
- Matchmaking.* SRs make a request represented by an SLA template to a broker or discovery service so as to find the SLA templates of those services that satisfy their quality requirements. This activity is after or in parallel with the functional service discovery activity. It may also be a joint activity of the two considered life-cycles, as SLAs can represent all information needed for the functional and quality-based service discovery. As a result of this activity, the user's SLA template may change, if it is overconstrained, to reflect realistic performance situations in the respective application domain.
- Negotiation.* This is a joint activity of the two considered life-cycles. The SP of the best service (or the SPs of the matched services) negotiates with the SR based on their SLA templates' content. These SLA templates may change during the negotiation to reflect the compromises made by the two parties.
- Agreement and Deployment.* As the outcome of service negotiation is not always successful, this activity is separated from the previous one. If the outcome is successful (agreement is reached), then a specific SLA is produced and signed by the two corresponding parties. This SLA has to be validated first and then deployed. SLA deployment is performed at two steps: (a) a signatory party's deployment system extracts from the SLA the appropriate configuration information and distributes it to the corresponding supporting parties so as to inform them about their roles and duties; (b) the supporting party deployment systems configure their own implementations in a suitable way. All parties need to know the definitions of the interfaces they must expose, as well as the interfaces of the partners they interact with. This composite activity is usually performed before service execution. When no active service instance can execute in the respective SL of the SLA, a new service instance must be deployed or more resources are given to a specific instance. In this case, the service deployment activity runs in parallel with the respective SLA activity.
- Monitoring and Assessment.* While the service is executing, the SLA is also "executed". The latter means that the service is periodically monitored and the SLA's agreed SLOs are assessed. Monitoring is performed by the measurement components

of the supporting parties' runtime systems which maintain information on the current system configuration and the SLA metrics. These components measure QoS metrics either from inside, by retrieving resource metrics directly from managed resources, or outside the SP's domain, for example, by probing or intercepting client invocations. The condition evaluation components of supporting parties' runtime systems compare the measurement information against the SLOs and notify the management systems (of the signatory parties and the SLA's external one (if it exists)) about violations. If the violation or violation number is very critical, then the SLA is renegotiated or canceled according to the SLA's corresponding action guarantee. If not, an appropriate corrective action is selected and performed by the SP's management system (e.g., more resources are provided to the underperforming service) according to the current context and the SP's business goals and policies. This SLA management activity runs in parallel with the service execution, monitoring, and recovery life-cycle activities.

- Settlement*. This activity determines if the SLA was met, the final cost to be paid by the SR, and which penalties may apply to the SP for breaching the SLA terms. Negotiations for SLA termination may be carried out between the parties, in the same way as the SLA establishment is performed, or for service reexecution in a different SL and cost. This activity occurs after the end of service execution.
- Archive*. After the SLA settlement takes place, the SLA may be disposed or archived according to the signatory parties' policies. However, even if the SLA is decided to be discarded, there is usually a statutory period (known as the "limitation period") where the SLA must be kept as it is a legal document describing how services were provided. If this activity is accompanied with an audit trail mechanism, it can be used for identifying problems and patterns of wrong service behavior or user requirement trends so as to improve the future developed SLA templates' content and even evolve the service implementation towards correcting the identified problems and meeting the increased customer needs. Thus, this activity can provide significant input to the service evolution life-cycle activity.

The summary of the selected criteria is shown in Table IX, while their complete presentation is provided later in this section. In this table, the term "Description" is used to cover the first two SLA life-cycle activities, that is, the "SLA Template Development" and the "SLA Advertisement", as they are actually associated with an SLA's description. Moreover, the "Agreement and Deployment" activity was not analyzed by using any criterion, since this activity does not need any specific information which is not already covered by existing SLA languages. Finally, as explained in Section 4, most SLA languages regard that quality should be defined outside the SLA specification by different formalisms and languages and only referenced inside this specification. For this reason, this section neglects from its analysis the quality description capabilities of the SLA languages.

Service contracts and SLAs are expressed by their respective languages which are shown in Table VIII. While there is no standard Service Contract Language (SCL), there are two widely used SLA languages, namely WSLA [Keller and Ludwig 2003] and WS-agreement [WS-Agreement 2003], which can be considered as standards.

The most representative SCLs have been chosen because, as was explained before, the majority of these languages focus more on the functional aspects of service behavior. The evaluation results of the examined languages according to our selected criteria are presented in Tables X and XI.

In the remaining part of this section, each activity-based group of criteria along with their evaluation results is presented in separate subsections. In the end, there is an overall analysis of the service contract and SLA languages across all criteria.

Table VIII. The SLA and Service Contract Languages Examined

ID	Approach Name	Approach Reference	Type
1	QML	[Frølund and Koistinen 1998]	SLA
2	WSLA	[Keller and Ludwig 2003]	SLA
3	WS-A	[WS-AGREEMENT 2003]	SLA
4	SLAng	[Lamanna et al. 2003]	SLA
5	WSOL	[Tosic et al. 2003b]	SLA
6	RBSLA	[Paschke 2005]	SLA
7	QoWL	[Brandic et al. 2006]	SLA
8	GXLA	[Tebbani and Aib 2006]	SLA
9	TrustCom	[TrustCoM Consortium 2007]	SLA
10	X-Contract	[Molina-Jimenez et al. 2003]	Service Contract
11	BCL	[Linington et al. 2004]	Service Contract
12	SweetDeal	[Grosf and Poon 2004]	Service Contract
13	CTXML	[Farrell et al. 2004]	Service Contract
14	SWCL	[Oren et al. 2005]	Service Contract

5.2.1. Description. Every SLA and service contract language must possess some properties that enable it to be a good candidate for representing SLAs. A set of such five properties/criteria has been selected and is analyzed in the next four subsections.

5.2.1.1. Formalism. Similarly to SQLs, each SLA language adopts a specific formalism to express its metamodel. These formalisms include: *informal* (such as DTDs or XML Schemas), *UML*, *RuleML*, *Finite State Machines* (FSMs), and *ontologies*, and have been used for categorizing each language.

The evaluation results are presented in the first row of the *Description* composite SLA life-cycle activity of Tables X and XI. The vast majority of the SLA-MMs are expressed with informal formalisms (mostly XML Schema but also other schema languages which focus on a language's concrete syntax). The same result applies to the SLA language partition but not to the SCL one, where SCLs exploit mostly formal formalisms. Schema languages are selected as they lead to a quick way of producing a language, surpassing in this way its abstract syntax expression. In addition, the majority of the approaches use XML Schema for the concrete syntax and XML for the SLA representation. XML is adopted due to its *platform independence*, *simplicity*, and *ease of use*, the excellent range of tool support available enabling *automatability*, including editors, parsers, transformation engines, and document validity checkers, and the fact that is both *human and machine understandable and processable*. However, XML Schema and correspondingly XML lack proper and precise *language semantics* needed to perform semantic SLA consistency and the *formality* needed to perform *SLA analysis* (analyzability) [Skene 2007] which can be used to reveal hidden obligations and other important nonvisible implications. For this reason, XML-based SLA descriptions are either transformed to another formalism [Linington et al. 2004] or other formalisms are adopted for expressing the SLA-MM like FSM (X-Contract), RuleML (RBSLA and SweetDeal), or Event Calculus (TCXML). This justifies the choice of SCL modelers to adopt such formalisms, as the focus is on analyzing the functional service behavior. While the adoption of stronger formalisms equips the SLA language with powerful tools to perform various forms of reasoning, there is usually a trade-off with simplicity, ease of use, and human understandability and processing.

Most SLA and service contract languages can perform at most the structural and semantical *validity* forms (i.e., discover syntactic and semantic inconsistencies with the help of DTDs, XML Schemas, and ontologies). However, another validity form is also required, called *SL validity*, in order to discover a specific *quality inconsistency*

Table IX. Summary of the Comparison Criteria of the SLA and Service Contract Languages

Life-Cycle Activity	Criteria	Summary
Description	Formalism	The language's description formalism
	Coverage	The ability to express functional and quality terms
	Reusability	The ability to represent SLA templates and other reusability constructs
	Composability	The ability to represent composite SLAs
Matchmaking	Metric Definition	The ability to define quality metrics
	Alternatives	The ability to express alternative SLs
	Soft Constraints	The ability to express <i>soft</i> SLOs that may be violated
	Matchmaking Metric	Definition of the way SLAs produced by a specific language must be compared
Negotiation	Meta-Negotiation	The ability to represent any information to be used for negotiation establishment
	Negotiability	The ability to define which SLA parts are negotiable and in what way
Monitoring	Metric Provider	The ability to define the party responsible of producing a metric's measurements
	Metric Schedule	The ability to define the production frequency of a metric's measurements
Assessment	Condition Evaluator	The ability to define the party responsible of SLO evaluation
	Qualifying Condition	The ability to define conditions that must hold in order to assess an SLO
	Obligated	The ability to express the party in charge of delivering what is promised in an SLO
	Assessment Schedule	The ability to express the assessment frequency of an SLO
	Validity Period	The ability to express the time period in which the SLO is guaranteed
	Recovery Actions	The ability to express corrective actions to be carried out when an SLO is violated
Settlement	Penalties	The ability to express penalties incurred when one party violates its promises
	Rewards	The ability to express rewards incurred when one party overwhelms its promise
	Settlement Actions	The ability to express actions concerning the final SLA outcome
Archive	Validity Period	The ability to express the period where an SLA is valid

type in SLs, which concerns the *constraint consistency*. As SLs are composed from the logical combinations of SLOs, it must be checked if this combination is meaningful and correct. For instance, if two constraints of the form $X < a$ and $X > b$, where X is a QoS metric and $b > a$, are conjunctively combined then metric X would not be allowed to take any value from its value type, so the produced SL would not be valid. This problem is exacerbated when arbitrary functions are involved in expressing SLO constraints. A good solution to this problem would be to transform the SL description into an appropriate constraint model and then use Constraint Logic Programming (CLP) (for arbitrary logical combinations of SLOs) or constraint programming (only for conjunctions of SLOs) techniques [Rossi et al. 2006] to check the constraint model's consistency [Müller et al. 2008].

5.2.1.2. Coverage. SLA and service contract languages should be able to express in an efficient and complete way both functional and quality terms. The functional term's description should include the description of the service functionality, operations, input,

Table X. Evaluation Results of SLA Languages

Life-cycle Activity	Criteria	QML	WSLA	WS-A	SLAng	WSOL	RBSLA	QoWL	GXLA	TrustCom
Description	Formalism	UML	inf.	inf.	inf.	inf.	RuleML	inf.	inf.	inf.
	Coverage	[p,y]	[p,y]	[y,p]	[y,y]	[p,p]	[p,y]	[y,p]	[p,p]	[y,y]
	Reusability	yes	yes	yes	part	yes	yes	part	yes	yes
	Composability	no	no	fair	fair	no	no	fair	good	fair
Matchmaking	Metric Definition	yes	yes	no	yes	no	yes	no	no	yes
	Alternatives	impl.	impl.	impl.	no	impl.	impl.	no	impl.	impl.
	Soft Constraints	no	no	yes	no	no	no	no	no	yes
	Matchmaking Metric	yes	no	no	yes	no	no	no	no	no
Negotiation	Meta-Negotiation	poor	poor	fair	poor	poor	poor	good	poor	fair
	Negotiability	no	no	part	no	no	no	no	no	part
Monitoring	Metric Provider	no	yes	no	yes	yes	no	no	no	yes
	Metric Schedule	no	yes	no	yes	no	yes	no	no	yes
Assessment	Condition Evaluator	no	yes	no	no	yes	no	no	no	no
	Qualifying Condition	no	yes	no	no	yes	no	no	no	no
	Obligated	no	yes	yes	yes	yes	yes	no	yes	yes
	Assessment Schedule	no	yes	no	no	no	no	no	yes	yes
	Validity Period	no	yes	no	no	no	yes	no	yes	yes
	Recovery Actions	no	yes	no	no	yes	yes	no	yes	no
Settlement	Penalties	no	no	SLO	no	SL	SL	SL	no	SLO
	Rewards	no	no	SLO	no	no	SL	no	no	SLO
	Settlement Actions	no	yes	no	no	no	yes	no	no	no
Archive	Validity Period	no	yes	yes	yes	no	no	no	yes	yes

and output. Moreover, if the service is composite, then all of its tasks, both internal and external, should also be described. All languages surveyed are able to define SLOs. However, the quality term's description should also include the description of the QoS metrics to be measured and various other quality concepts in order to be complete. If the SLA or service contract languages are not able to describe these additional quality terms, they reference these terms' external descriptions in respective QSLs. Thus, this criterion's evaluation considers these two description aspects and provides the following values for a language: $[n, n]$ (no functional and no SLO description), $[n, p]$ (no functional and only SLO description), $[n, y]$ (only complete quality description), $[p, p]$ (references to functional and only SLO description), $[p, y]$ (reference to functional and complete quality description), $[y, y]$ (complete functional and quality description).

The evaluation results are presented in the second row of the *Description* activity of Tables X and XI. All the languages reference or explicitly define functional terms. Besides, the majority of these languages reference an external functional description of the offered service (e.g., WSDL or BPEL). This result also applies in each partition. Only four SLA languages enable the description of the functional terms of the SLA, while no SCL enables this description. Moreover, the WS-agreement and TrustCom SLA languages enable both the description and reference of the functional SLA terms.

Table XI. Evaluation Results of Service Contract Languages

Life-cycle Activity	Criteria	X-Contract	BCL	SweetDeal	CTXML	SWCL
Description	Formalism	FSM	inf.	RuleML	inf.	ontology
	Coverage	[p,n]	[p,n]	[p,n]	[p,p]	[p,p]
	Reusability	no	yes	yes	part	part
	Composability	no	no	neutral	no	no
Matchmaking	Metric Definition	no	no	no	no	no
	Alternatives	impl.	no	no	no	no
	Soft Constraints	no	no	yes	no	no
	Matchmaking Metric	no	no	no	no	no
Negotiation	Meta-Negotiation	fair	poor	poor	poor	poor
	Negotiability	no	no	no	no	no
Monitoring	Metric Provider	no	no	no	no	yes
	Metric Schedule	no	no	no	no	no
Assessment	Condition Evaluator	yes	no	yes	no	no
	Qualifying Condition	no	no	no	yes	yes
	Obligated	yes	yes	yes	yes	yes
	Assessment Schedule	yes	yes	yes	yes	yes
	Validity Period	no	yes	no	no	no
	Recovery Actions	yes	yes	yes	yes	yes
Settlement	Penalties	no	SLO	SL	SL	SL
	Rewards	no	no	SL	SL	SL
	Settlement Actions	yes	yes	yes	yes	yes
Archive	Validity Period	no	yes	no	no	no

The latter language is an WSLA-based extension of WS-agreement, so it is reasonable to inherit the majority of WS-agreement's capabilities. As already explained in Section 4, the syntactical separation of functional and SLA descriptions enables the reuse of a specific contract among many services that exhibit the same quality capabilities. In addition, it facilitates the contract management and evolution (through the manipulation of changing SLs [Tosic et al. 2003b]) and disables the repetition of a service's functional description in all the service's contracts. On the other hand, the inclusion of the service functional description inside an SLA mitigates the risk that the client experiences a different functionality from the one requested without being able to claim for this violation in the agreement. In this way, abnormal behavior of SPs that change the service functionality and description externally and invisibly to an SLA is avoided.

Concerning the description of additional quality terms apart from SLOs, almost half of the languages only reference external quality descriptions usually found in SQMs. In the SLA language partition, there is no language that does not describe or reference additional quality terms and this result is quite reasonable. The CSLs that have been constructed to accommodate for any possible electronic contract and not just SLAs are not very efficient in this matter, as most of them do not reference any quality term but could be extended to do so. It must be noted that referencing external descriptions of the additional quality terms enables their reuse but may create problems in the *match-making* activity of the SLA management life-cycle, which is analyzed in the next item.

5.2.1.3. Reusability and extensibility. An SLA/service contract language should enable the creation of templates and documents that can be reused or extended for creating new SLA or service contract specifications, respectively. Moreover, SLA or service contract specification parts, like the functional and quality terms, should be reused across many SLA or service contract documents or extended appropriately. Thus, this

criterion's evaluation for each language could get the following values: *no*, *part* (i.e., only the whole SLA/service contract is reusable and extensible), and *yes* (i.e., parts and whole SLA/service contract are extensible and reusable).

The evaluation results are presented in the third row of the *Description* activity of Tables X and XI. Most SLA languages are reusable and extensible both in the SLA specification entirety and in its parts. This result applies also to the SLA language partition. Concerning the SCL partition, there is a balance between those SCLs that are reusable and extensible only in their entirety and those that are also reusable and extensible in their parts.

5.2.1.4. Composability. This property indicates the ability of an SLA or service contract language to represent composite SLAs or service contracts, respectively. It can be achieved when the language presents the following abilities: (a) describe or reference composite service descriptions; (b) define or reference metrics that are associated to composite services and are computed based on aggregation rules that depend on the composite service structure; (c) cater for the different party types (i.e., third-party SPs) involved in composite contracts; (d) define composite and component-based SLs and their associations; (e) define appropriate action guarantees that consider the contract's two-level hierarchy. Most of the languages possess at most both of the first two abilities. For this reason, this criterion has been evaluated for each language depending on the language's satisfaction of these five abilities according to their order. So, if the language does not possess any ability, it is evaluated with *no* (i.e., it is not composable). If it possesses one of the first two abilities, it is *neutral*. If it possesses the first two abilities, it is considered *fair* (i.e., it has the basis for becoming composable). If it possesses the first three or four abilities, it is considered *good*. Finally, if it possesses all the abilities, it is evaluated with *yes* (i.e., it is composable).

The evaluation results are presented in the last row of the *Description* activity of Tables X and XI. Most languages have not even the basis of being composable. This result applies also to the SCL partition, where four out of five approaches are not composable at all, while one approach possesses only the first ability. This actually means that SCLs were not designed to represent composite service contracts or SLAs. Concerning the SLA languages' partition, there is a balance between approaches that have fair/basic and no composability at all. Only one SLA (GXLA) language scores "good" in its composability. This means that the SLA modelers have not yet understood the need for representing composite SLAs.

5.2.1.5. Overall analysis for the description activity. Based on the preceding analysis, there is no SLA or SCL that meets the high standards posed. Only GXLA can be distinguished based on its capabilities with respect to composability. However, this language lacks the appropriate formality (or transformation to such a formality) which is needed for SLA validation and analysis.

5.2.2. Matchmaking. Only one research approach [Oldham et al. 2006] performs proper matchmaking of SLA specifications so as to match the user QoS requirements with the service QoS capabilities. In this approach, WS-agreement specifications are enriched with semantic annotations from both domain-independent and domain-dependent ontologies, whilst rules are also used to infer the matchmaking. The usual procedure followed in the remaining approaches is that matchmaking is performed during service negotiation, where one participant proposes a specific SLA (or service contract) template and the other one accepts it or changes it, by implicitly checking every SLO and changing its limits and by entering new SLOs, or proposes a new one. This matchmaking type is inefficient for the following reasons: (a) QoS metrics are defined inadequately and syntactically in possibly different languages leading to low-accuracy

results; (b) no matchmaking metric is defined so each party utilizes its own metric to infer if the proposed SLA template is the appropriate one; (c) the probability of matching is low because SLOs are usually expressed as hard constraints; (d) it is time consuming as in each iteration the receiving party has to parse and process an SLA template, compute the answer, and send a modified version or a new template.

Based on the content of most SLA life-cycles (including the one proposed) and the previous reasoning, SLA/service contract template matchmaking should be performed before service negotiation in order to discover those services of the corresponding SPs that suit the user quality requirements. This process can be effective and accurate if some prerequisites are met by the SLA and service contract languages and there is a common, unique, and fair matchmaking metric that can be used to perform the matchmaking in such way that will always yield the same results for the same input. In the following, three main description prerequisites and one specific requirement (i.e., the existence of a matchmaking metric) are analyzed, which must be met by an SLA language in order to enable the matching of its specifications. In the end, an overall analysis of the ability of the SLA languages to support this SLA management life-cycle activity is given.

5.2.2.1. Metric definition. Metric modelling capabilities were analyzed in Section 4, where four SLA languages were compared. The rest of the SLA (apart from TrustCom) and service contract languages considered are not able to define QoS metrics but just reference external metric descriptions of SSQLs. In this way, two main problems may arise: (a) language incompatibility may occur, wherein the involved SSQLs may encompass different metric metamodels so they can describe metrics in a different way and, thus, it will be difficult to transform one SSQL's metric description into the other's one when matchmaking SLA descriptions; (b) even if the SSQLs are compatible, equivalent metrics described in these different languages may have a different name, so their descriptions have to be matched via metric matching rules [Kritikos and Plexousakis 2006] to infer their equivalence. These two problems reduce the matchmaking activity's accuracy. Thus, languages that define metrics or enforce the use of a specific SSQL to define metrics are preferred. So, this criterion's evaluation for each language would be *no*, if the language does not define metrics or reference metric descriptions from a specific SSQL, or *yes* otherwise.

The evaluation results are presented in the first row of the *Matchmaking* life-cycle activity of Tables X and XI. Five SLA languages (four plus TrustCom that relies on WSLA) satisfy this criterion, which are the approaches able to define QoS metrics and other SLO quality terms on their own. So, by also considering the results of the *coverage* criterion of the *Description* composite activity, it can be inferred that all languages that can reference external metric descriptions do not determine which SSQL should be used to specify these descriptions. Thus, based on the preceding analysis, the use of these languages may lead to the low accuracy of the matchmaking results.

5.2.2.2. Alternatives. Two SLA/service contract specifications match when their part corresponding to the offered or required SL is matched. Existing languages do not encompass the SL concept but either assume it, is the conjunction of all SLOs defined in the SLA or emulate it, either through the use of logical predicates that logically connect the defined SLOs or by offering different SLAs for each service. However, the ability to implicitly represent an SL is not enough as the probability that there is a match between the encompassing SLOs of the compared SLs of two SLA specifications is very low. Moreover, users have diverse needs that can be represented through trade-offs between the requested SL and its cost. Thus, SLA and service contract languages should be able to represent alternative SLs in order to increase the chances of matching their corresponding specifications. Alternative SLs represent the different modes in

which a service can operate to suit the diverse needs of different user classes, and the variations of a requested SL by an SR that trade off the SL with the price the SR is willing to pay. Thus, this criterion's evaluation for each language would be *no* if the language is not able to represent SLs, *impl* if it can define them implicitly, and *yes* if there is an explicit language construct that is used to represent alternative SLs.

The evaluation results are presented in the second row of the *Matchmaking* life-cycle activity of Tables X and XI. Most language's can implicitly define alternative SLs which are needed to increase the chances of matchmaking with potential SRs. This result also applies to the SLA languages' partition. Moreover, there is no language that explicitly defines SLs. Finally, among the SCLs, only X-Contract can implicitly define SLs. Thus, both SLA and service contract languages were not designed to support SLs but only SLA languages can implicitly define them.

5.2.2.3. Soft constraints. Even if many alternative SLs are represented in an SLA template to be matched, there will always exist a problem [Kritikos 2008] where users express overconstrained SLs that cannot be matched by any alternative SL of any offered SLA template. An overconstrained SL means that its encompassing SLOs contain very restrictive constraints that cannot be satisfied. As the root of this problem is that SLOs are expressed as hard constraints that must be satisfied at all costs to infer the matchmaking, its solution may come through the use of soft constraints [Zemni et al. 2010]. In particular, if SLOs are expressed as soft constraints, where the user expresses their significance through using weights or levels, then not all of them have to be satisfied when matching. In this way, there can be a match between an offered and requested SL, even if some insignificant requested SLOs are violated. Thus, this criterion's evaluation for each language would be *no* if the language cannot express soft SLO constraints, or *yes* otherwise.

The evaluation results are presented in the third row of the *Matchmaking* life-cycle activity of Tables X and XI. Only two SLA and one service contract language can define soft constraints, while the rest define SLOs as hard constraints. Thus, only these three languages could be used to express SLAs/service contracts that could be exploited to solve the overconstrained user-requested SLs problem.

5.2.2.4. Matchmaking metric. As languages may differ in the way they define QoS metrics and SLOs, it would be very useful when implementing SLA matchmaking engines if a specific matchmaking metric was defined internally or externally in the SLA language. This metric would be used for matching SLAs (defined by the respective language) in a fair and uniform manner according to the matchmaking requirements defined in Kritikos [2008] and Kritikos and Plexousakis [2009]. Thus, this criterion's evaluation for each language would be *no* if no matchmaking metric is defined, or *yes* otherwise.

The evaluation results are presented in the last row of the *Matchmaking* life-cycle activity of Tables X and XI. Only two SLA languages (QML and SLang) explicitly define a matchmaking metric with which their SLA specifications can be matched.

5.2.2.5. Overall analysis for the matchmaking activity. There is no language that satisfies all four criteria of the *matchmaking* activity. Only QML and TrustCom satisfy three of the criteria. However, among these two languages QML is considered as the best because the *soft constraints* criterion is the least significant one as it provides additional and not basic support the SLA matchmaking activity. QML is the oldest of all SLA languages and is not used any more. However, it was designed with the explicit goal of *contract conformance*, which is actually used for the SLA specification matchmaking. Thus, it can be deduced that the majority of the languages and especially

the SCLs were not designed with the objective of matchmaking the quality terms of their specifications.

5.2.3. Negotiation. Service negotiation is one of the most important activities as it produces the final SLA document that will drive the service execution. For this reason, SLA languages must describe all the appropriate information that will be provided as input and assist the service negotiation process. This information can be categorized into two parts: *meta-negotiation*, and *negotiability*. In the following, it is explained why these two information types must be captured by SLA languages and what should be their content, while it is inspected if the languages capture this essential information. Finally, the overall performance and support of the examined languages for this SLA management activity are assessed. As *negotiation strategies* represent private and sensitive information for the participants which must not be exposed in SLA templates, it was considered that they do not constitute appropriate SLA information for the support of service negotiation.

5.2.3.1. MetaNegotiation. *Metanegotiation* is any information that can be used for negotiation establishment, that is, to enable and initiate the negotiation between the participants. The following information has been identified as metanegotiation [Brandic et al. 2009; Comuzzi et al. 2009]: (a) negotiation protocol support; (b) description of negotiation capabilities; (c) authentication method reference.

The *negotiation protocol* is the allowable sequence of exchanged messages used to negotiate and conclude (i.e., agree) an SLA/service contract. The protocol should also unambiguously define the semantics and format, or schema, of the messages. Each negotiation participant may be able to support a subset of all possible negotiation protocols. Thus, the supported negotiation protocols of all participants must be matched in order to find the one appropriate for enacting the negotiation. For this reason, each participant's corresponding SLA/service contract template must reference all negotiation protocols that can be supported. Moreover, this reference should include a pointer to the supported negotiation protocol description for two reasons: (a) to enable reasoning on protocol compatibility and substitutability, that is, if one protocol can be used in place of the other, and (b) some negotiation protocols may not be implemented in negotiation engines and brokers, for example, are not widely used; therefore, they have to be defined, for example, in BPEL, in order to be properly enacted by the negotiation broker.

When no matching negotiation protocol is found, it is advocated in Comuzzi et al. [2009] that the participant negotiation capabilities must be described in a more fine-grained way along with the possibilities of delegating capabilities to trusted third parties. Then, ontology-based reasoning can be used to infer if a specific negotiation protocol can be supported by the participants and their trusted third parties. Thus, based on the earlier analysis, the participant's fine-grained negotiation capabilities must be also advertised in their SLA templates apart from the coarse-grained ones.

While negotiating, the participants reveal and exchange important information which should not be exposed to third parties listening on the insecure channels established between the participants. For example, SPs do not want the offers they make to specific (e.g., privileged) clients to be viewed by all other clients. For this reason, the participants must describe in their SLA templates the authentication methods they prefer to be used for securing the channels exploited when exchanging negotiation information. Then, preferred authentication methods will be matched so as to select the most common one for the negotiation.

Each language is evaluated according to the number of metanegotiation information it can describe. Thus, if it cannot describe any information, it is evaluated as *poor*. If it describes only one of the aforesaid information types, it is considered *fair*. If it describes

two metanegotiation information types, it is considered *good*. Finally, if it describes all possible negotiation information, it is considered *rich*.

The evaluation results are presented in the first row of the *Negotiation* life-cycle activity of Tables X and XI. The vast majority of the languages are not capable of modelling any metanegotiation information. This result holds also in every partition. Only WS-agreement, TrustCom, and X-Contract determine or can represent negotiation protocols, while QoWL is able to both specify the negotiation protocol and the authentication method to be used for the SLA negotiation.

5.2.3.2. Negotiability. While *meta-negotiation* represents information that is used to support the service negotiation enactment, *negotiability* represents information which is used during the negotiation. In particular, negotiability is the ability of an SLA language to describe which parts of its specifications are negotiable and in what way. Focusing on quality, a language presents negotiability if it can characterize which quality terms are negotiable or not and which are the allowable values in the quality term's upper and lower limits. Thus, this criterion's evaluation for each language would be *no* if the language does not define which terms are negotiable, *part* if the language characterizes only the term's negotiability, and *yes* if the language also determines which are the allowed values or range of values for the quality term's upper and/or lower limits.

The evaluation results are presented in the second row of the *Negotiation* life-cycle activity of Tables X and XI. Only two SLA languages (WS-agreement and TrustCom) specify in a special part of their produced SLA templates which terms are negotiable. However, they do not specify the way these terms are negotiable. On the contrary, the constraints used to define the negotiable terms are hard and do not specify if a quality term's limits can take one or more values.

5.2.3.3. Overall analysis for the negotiation activity. According to the evaluation results of the two criteria of the *negotiation* activity, only WS-agreement (and TrustCom that extends it) can partially provide information that can assist this activity. However, they should be extended by modelling the participant negotiation capabilities, the authentication method used for the information exchange during the negotiation, and specific constraints that define the allowed values for the negotiable quality term's limits. Considering the latter extension, one good solution is proposed in Andrieux et al. [2004] which is, however, not adopted in the language's formal specification. SCLs are not able to properly support the quality-based term's negotiation. This fact along with the inability of the SCLs to support the SLA matchmaking activity prevents them from being widely adopted as the languages for expressing a service's quality-based behavior.

5.2.4. Monitoring and Assessment. During service execution, the service is monitored so as to assess if the SLOs defined in its SLA are violated. Service monitoring is performed by producing measurements according to the information that is encapsulated in the SLO metrics that are in the service scope. As both the metric definition and the association of metrics to service objects were evaluated in previous parts of this article, the only additional information needed for *monitoring* is who is in charge of performing the metric measurements, that is, the *Metric Provider*, and how often the measurements are produced, that is, the *Metric Schedule*. This information is encompassed in a metric metamodel, but is usually specified concretely only when the SLAs are established after service negotiation.

5.2.4.1. Metric provider. The Metric Provider is the responsible party for producing a specific metric's measurements. This criterion's evaluation for each language would be *no* if the language does not define this party, or *yes* otherwise.

The evaluation results are presented in the first row of the *Monitoring* life-cycle activity of Tables X and XI. Almost half SLA and almost all service contract languages are not able to specify the providers of SLO metrics as they assume that the measurements are only provided by the SP. This is a major limitation because measurements may be provided by other parties and the SLO evaluators, which may be different from SPs or even SRs, would not be able to assess the SLOs if they do not know the place from where SLO metric measurements can be obtained.

5.2.4.2. Metric schedule. The schedule of a metric determines the production frequency of its measurements. This criterion's evaluation for each language would be *no* if the language cannot define metric schedules, or *yes* otherwise.

The evaluation results are presented in the second row of the *Monitoring* life-cycle activity of Tables X and XI. Only four SLA languages are able to specify metric schedules. Three (WSML, WSLA, and SLAng) of these four languages satisfy both the current and the previous criterion. The rest of the SLA languages and all SCLs do not model this feature. However, as this feature is used to specify the timing of the measurement productions of the SLO metrics, its lack can cause problems in the *assessment* activity.

The SLA assessment is one of the most critical SLA life-cycle activities. Thus, SLA languages must model all appropriate information that could be used to support this activity. Apart from the main condition of the SLO (clause) that is modelled in all SLA/service contract languages, other important SLA assessment information that should be modelled is: (a) *Condition Evaluator*, (b) *Qualifying Condition*, (c) *Obligated Party*, (d) *Assessment Schedule*, (e) *Validity Period*, and (f) *Corrective Actions*. In the following, the purpose and content of this information is analyzed and it is evaluated if the examined languages have modelled it.

5.2.4.3. Condition evaluator. Similarly to metric measurement, the SLO assessment should be made by a (supporting) party which is named *Condition Evaluator*. This party is in charge of collecting the measured values of all metrics involved in an SLO, replacing the metrics with their values, and then checking if the SLO holds or not. This criterion's evaluation for each language would be *no* if the language does not define this type of supporting party, or *yes* otherwise.

The evaluation results are presented in the first row of the *Assessment* life-cycle activity of Tables X and XI. Only two SLA languages (WSLA and WSOL) and two SCLs (X-Contract and SweetDeal) are able to model this information. All the other languages presuppose that the SLA assessment activity is performed in the signatory party's management systems based on the information coming or pulled from the monitoring components. In this way, they limit the way an SLA management system can be implemented or distributed as they exclude the existence of third-party assessment components.

5.2.4.4. Qualifying condition. Apart from the SLO to be evaluated, there can be a precondition, the *qualifying condition*, which must hold to assess the SLO. This precondition may express assertions over service or other quality attributes or external factors such as the SR's *service request rate*. This criterion's evaluation for each language would be *no* if the language does not define qualifying conditions, *impl* if it defines them implicitly through other constructs, or *yes* otherwise.

The evaluation results are presented in the second row of the *Assessment* life-cycle activity of Tables X and XI. Only one SLA language (WS-agreement) and two SCLs are able to explicitly model the *qualifying condition* attribute, while another SLA language (WSLA) can implicitly define it through other constructs. All the other languages do not offer this capability. This lack leads to inability of expressing preconditions for the enactment of an SLO's assessment, which would eventually lead to situations where

SLOs are assessed wrongly with regards to timing or other excluding conditions (e.g., client-side or management-related restrictions).

5.2.4.5. Obligated. The *Obligated* party is in charge of delivering what is promised in an SLO. In many cases, this party is the SP, while in other cases it can be another party, such as a service component's third-party provider. Thus, every language should associate an SLO with the party that promises it. This criterion's evaluation for each language would be *no* if the language does not define the obligated party in the SLOs, or *yes* otherwise.

The evaluation results are presented in the third row of the *Assessment* life-cycle activity of Tables X and XI. Most SLA languages are able to define which is the obligated party in an SLO. Moreover, all SCLs model this information as their design is based on policies expressing obligations and various other implication types.

5.2.4.6. Assessment schedule. An SLO is not assessed just one but several times according to an *Assessment Schedule*. This schedule can be as simple as assessing when new values are measured for the SLO metric or complex, representing a sequence of regularly occurring events. This criterion's evaluation for each language would be *no* if the language cannot define assessment schedules, or *yes* otherwise.

The evaluation results are presented in the fourth row of the *Assessment* life-cycle activity of Tables X and XI. Most of the languages specify an SLO's assessment schedule. This result applies also to the SCL partition, as SCLs have been designed to support this criterion. However, most of the SLA languages do not model this information. This limits their application only in expressing SLAs that involve services whose performance should be checked at only one instant. The languages that model this criterion follow two different approaches. In the first approach the schedule is defined concretely with timing constraints, while in the second approach the schedule is based on events originating from the SLA management system's monitoring components. The first approach is adopted by SLA languages, while the second approach is adopted by all SCLs.

5.2.4.7. Validity period. While the assessment schedule determines when to assess an SLO, the validity period determines the time period in which the SLO is guaranteed and, thus, should be checked for validity. An example of the value set this field can take is {*business days, regular working hours, maintenance periods*}. This criterion's evaluation for each language would be *no* if the language does not define the an SLO's validity period, or *yes* otherwise.

The evaluation results are presented in the fifth row of the *Assessment* life-cycle activity of Tables X and XI. Most of the language's do not model this information. This result also applies to the SCL partition. However, it does not apply to the SLA languages' partition, as there is a balance between the approaches that model this information and those that do not. Thus, SLA language designers have better understood the need of supporting this criterion with respect to those of the SCLs.

5.2.4.8. Corrective actions. When the signatory parties are informed about an SLO violation, corrective actions must be carried out at the obligated party's management system or at the global level by renegotiating or canceling the SLA/service contract. When corrective actions should be taken by the obligated party, the choice of which action to perform depends on many situational factors and the obligated party's business goals and policies. So, this party would not desire or be possible to advertise in an SLA/service contract what actions to perform in which SLO violation case. However, this party may advertise some corrective actions to be taken for the corresponding SLO's violations to increase its reputation and trust with respect to the SR or to guarantee the high gain that will come from assuring the agreed SL to a *golden class* customer. Thus,

this information should be certainly modelled in an SLA/service contract. So, this criterion's evaluation for each language would be *no* if the language does not model *corrective actions*, or *yes* otherwise.

The evaluation results are presented in the fifth row of the *Assessment* life-cycle activity of Tables X and XI. The majority of the language designers have recognized the significance of modelling this information. The same result also applies to the SCL partition. However, it does not apply to the other partition, as there is a balance between SLA languages that model corrective actions and those that do not. This means that SCL designers have better understood the importance of this criterion with respect to those of the SLA languages.

5.2.4.9. Overall analysis for the monitoring and assessment activity. Based on the overall performance of the examined languages on the monitoring and assessment criteria, WSLA is the best language that models all appropriate information that is required for supporting the *Monitoring* and *Assessment* life-cycle activities, as the support of these two activities was one of its design requirements and seems to be successfully implemented. By inspecting the two different partitions, SCLs offer well support, while apart from WSLA the rest of the SLA languages do not seem to support well the SLA monitoring and assessment activities.

5.2.5. Settlement. This activity assesses what has happened during the service's execution and what are each signatory party's responsibilities according to the agreed SLA. So, for example, if a specific SLO was violated, then the SP has to pay a small penalty to the SR. As another example, if the service runs in a higher SL than requested, then the SR has to pay, apart from the actual service cost, a reward for getting a better SL. The appropriate information to be modelled by an SLA/service contract language for supporting this activity is the following: (a) the incurred penalties, (b) the incurred rewards, (c) settlement actions.

5.2.5.1. Penalties. Penalties are paid by the SP if one or more SLOs are violated. Each SLO is usually associated with a specific penalty amount. However, in some cases, the penalty to be paid could increase exponentially with the violation number [Paschke and Schnappinger-Gerull 2006]. The latter penalty type is not modelled by most languages as it would require the definition of the appropriate SL first and then its association to a specific policy or function that would increase exponentially or linearly according to the violation number in the SLOs that compose this SL. If a language is able to define penalties at the SL, then it can also define penalties at the individual SLO level. Thus, this criterion's evaluation for each language would be *no* if the language does not define penalties, *SLO* if it defines penalties at the SLO level, or *SL* if it defines total penalties at the SL level.

The evaluation results are presented in the first row of the *Settlement* life-cycle activity of Tables X and XI. Most of the languages are able to specify penalties. The majority of these languages can define penalties for the whole SL, while only three languages define penalties for each SLO. Considering each partition separately, the majority of the SCLs are able to define penalties for each SL, while there is a balance between those SLA languages that are able to define penalties and those that are not. Moreover, there is a balance between the SLA languages that model penalties at the SL level and those that model penalties at the SLO level. The results show that SCL designers have better understood the need to model penalties, as opposed to the SLA language designers.

5.2.5.2. Rewards. Rewards are paid by the SR if one or more SLOs are more than respected. Rewards can be defined at the SLO (via a specific value) or SL level (via a function). Rewards should be modelled as they would give extra motives to SPs to

provide even better SLs than the ones offered to respective SRs in the past. This SL upgrade would lead to increase in profits, which is one main goal of SPs when they offer their services. Moreover, the SL's trust and reliability would also increase. This criterion's evaluation for each language would be *no* if the language does not define rewards, *SLO* if it defines rewards at the SLO level, or *SL* if it defines rewards at the SL level.

The evaluation results are presented in the second row of the *Settlement* life-cycle activity of Tables X and XI. The results are different with respect to those of the previous criterion. Less than half of all languages are able to define rewards. Moreover, there is a balance between the languages able to define rewards at the SL level and those able to define rewards at the SLO level. In the SCL language partition, all SCLs are able to define rewards at the SL level. Concerning the other partition, the majority of the SLA languages are not able to define rewards as either no real-world case of SLAs contains rewards or the SLA language designers have not recognized the need of modelling rewards. As SLA languages represent the majority of all languages, this explains the bad global result.

5.2.5.3. Settlement actions. Settlement actions are mutually taken by both signatory parties to decide about the SLA/service contract final outcome. Thus, when there are no severe SLO violations or the violation number is not high or zero, the SLA outcome is successful and maybe only penalties or rewards have to be paid. However, in the opposite case, it should be determined if the SLA must be canceled, renegotiated, or reenforced (e.g., the service has to be reexecuted). Thus, SLA/service contract languages should be able to model these settlement actions and the conditions on which they are applied. So, this criterion's evaluation for each language would be *no* if no settlement actions can be defined, or *yes* otherwise.

The evaluation results are presented in the last row of the *Settlement* life-cycle activity of Tables X and XI. Half of the languages are not able to model settlement actions. Concerning the SLA partition, only two SLA languages are able to model such actions. This is a significant limitation that would discourage potential SPs or SRs from using them. Moreover, this explains the bad global result. This situation is reversed in the SCL partition, as all SCLs support the modelling of this information. This means that, indeed, the SCL design has been centered on the modelling of various compensation actions, including the settlement ones.

5.2.5.4. Overall analysis for the settlement activity. Based on the evaluation results of the settlement criteria, only RBSLA among the SLA languages satisfies all criteria and can specify penalties and rewards on the SL. The same applies for three (SweetDeal, CTXML, and SWCL) out of five SCLs. Thus, these four languages can be used for appropriately supporting the *Settlement* life-cycle activity. However, these languages require significant effort and extensions from the SLA modeler to express different SLAs with different settlement actions and conditions. In addition, they force the SLA management systems adopting them either to support one by one the different settlement actions that may exist in the SLAs or to define appropriate extensions with which the modelers may specify their SLAs. The latter may lead to a situation where various different versions of the same language are adopted by different SLA/service contract management systems. Thus, to avoid such situations and further increase their adoption and universality, these languages should be extended appropriately to specify explicit constructs that model the various settlement actions that may exist in an SLA/service contract.

5.2.6. Archive. An SLA/service contract is archived in three distinct cases: (a) the SLA is canceled, (b) the maximum number of service invocations has been reached, or (c) its

validity period has expired. In the first case, settlement or corrective actions dictate when the SLA is canceled. In the second case, the SLA has to determine this maximum number of service invocations, and none of the existing languages is able to model this information. In the third case, the language has to model the SLA's validity period. Besides the timing of SLA archiving, some parties desire to dispose of the SLA. In this case, the SLA is first archived and then disposed of when a specific statutory period is expired. Again, none of the existing languages models this information. Thus, for this activity, each language is evaluated only based on its capability to model the SLA validity period. If it does not model this period, the evaluation result is *no*, otherwise *yes*.

The evaluation results are presented in the last row of Tables X and XI. While the modelling of an SLA's (service contract's) validity period is very important, less than half of the languages are able to offer it. Concerning each partition separately, there is a balance between SLA languages that support and do not support the modelling of this information. On the other hand, only one SCL models this information. So, the SLA language designers are starting to understand the need of modelling this information, while the corresponding SCL designers do not. From the SCL side, this conclusion can be explained by considering that these languages were designed with the focus on functionality and not on quality. Thus, as functionality does not change so much, there is no need to explicitly model the service contract validity period. The contract could be invalidated as soon as the business relationship between the contracting organizations ceases to exist for various reasons.

5.3. Overall Analysis

The analysis performed has revealed some significant facts and limitations of existing SLA and service contract languages in their ability to appropriately support the SLA life-cycle management activities. In this subsection, the analysis focuses on the overall, global level of SLA management activity support so as to reveal other interesting facts that are not obvious in a first sight.

First of all, by inspecting each activity's overall results, it can be inferred that there is no language supporting in a satisfactory way all activities. On the contrary, in each activity a different language is awarded as the most appropriate one. In addition, there are some SLA management activities that are not satisfactorily supported by any language, including those of SLA Description, Matchmaking, and Negotiation, while others are properly supported by few languages, including those of SLA Monitoring and Assessment, Settlement, and Archive.

The preceding general results burden the SLA languages that were explicitly designed to support all SLA management activities, as they signify that these languages do not possess all appropriate modelling capabilities so they should be used for expressing only some SLA types. Based on the fact that the capabilities of the most widely used SLA languages, that is, WSLA and WS-agreement, are complementary with respect to the SLA management activities support, one solution that could be adopted is to design a new SLA language that unifies the capabilities of these languages by extending them and encompassing some modelling constructs of the one to the other. One such paradigm is TrustCom, which is the only SLA language that has a good evaluation score across all the activities. Another solution would be to design a new SLA language that could use the best modelling features of the two standardized ones and explicitly model the missing features.

Another interesting result derived from the previous analysis is that SCLs do not fully support most of the SLA management activities apart from those of SLA Monitoring and Assessment and Settlement. This is due to the focus of SCL design on service functionality, which was inevitable during SCL modelling time. In this way, service quality, which has the main focus now because of its dynamicity, is either neglected or

not appropriately modelled. Thus, although these languages were designed to accommodate for any electronic contract type, they cannot be used for specifying SLAs unless they are extended appropriately. SWCL could be easily extended as it has the best score among all SCLs across all SLA management activities. This language along with TCXML are the most representative and recent SCLs which have included quality-related constructs in order to accommodate for the change of focus from service functionality to quality.

6. CONCLUDING REMARKS AND FUTURE WORK

This article has focused on investigating the issue of service quality description. To this end, a systematic review of a large number of approaches has been conducted in order to reveal their strengths and weaknesses and highlight where the need for further research and investigation belongs. Initially, the approaches were separated into three clusters according to their scope: (1) *service quality models* (SQMs) which are taxonomies of service quality that can be used to annotate other types of quality documents like QSDs and SLAs, (2) *service quality metamodels* (SQMMs) which are capable of expressing SQMs as well as service quality offerings and requirements (i.e., QSDs), and (3) *service-level agreement metamodels* (SLA-MMs) which are capable of describing SLAs. Then, there was a comparison of the approaches of each cluster according to a set of scope-specific criteria aiming at unveiling which approaches are the consolidated ones and which are the ones specific to given aspects. This comparison uncovered many interesting findings and has also identified particular aspects of underperformance concerning each cluster's approaches.

The next three subsections summarize the most important of these findings and draw directions for further research and improvement.

6.1. Discussion on Service Quality Models

Various SQMs have been proposed, from small or flat categories of service quality attributes to sophisticated taxonomies containing many categories and attribute types. In order to compare these approaches in a fair and consistent manner, a set of criteria were devised characterizing the extensiveness, information richness, structure, generality, and applicability of the considered SQMs. Concerning the first four aspects, the evaluation results have shown a trend that the approaches are improving over the years. In average, the SQMs have a satisfactory category number, where each category contains a small quality attribute number. Most SQMs mainly cover general (i.e., domain-independent) quality attributes, while a small number of them also cover specific (i.e., domain-dependent) ones. As the inclusion of general attributes tends to cover the SP view while the inclusion of specific ones tends to cover the SR view, most SQMs mainly cover the SP view. Besides, most SQMs contain both composite and atomic quality attributes along with the connecting relation between them. The latter relation is very important during service monitoring as it may be used to validate or enrich the monitoring results of a service monitoring engine or component.

Another interesting finding is that the majority of the SQMs include only QoS attributes but only the most recent approaches also include QoE attributes. The latter result demonstrates how researchers in this field are starting to realize that quality attributes that are assessed subjectively from user feedback are as important as objectively assessed quality attributes. This is because subjective attributes reveal the service's performance and usability from the perspective of actual service users, so they also constitute critical service selection factors.

An important evaluation result shows that, compared to the very first approaches which focused on quality attributes in the service layer, current approaches are increasingly considering attributes that can be associated to one or both of the two other

layers in our quality model, namely the application and infrastructure layers. While this is a significant advance, it is outweighed by the lack of interattribute dependencies not only within the same but also across the layers. Interattribute dependencies are extremely important as they reveal the influence one attribute has on the other. In this way, the service monitoring, assessment, and adaptation activities can exploit them to perform dependency analysis in order to detect wrong monitoring facts and discover the components of the same or different layer that caused an SLO violation.

The current state-of-the-art approaches scarcely consider data quality aspects. However, since the service output is mostly composed of information, data quality can be considered as a part of service QoS and can drive thoroughly the analysis of the required input and provided output.

Concerning the applicability comparison aspect, only one SQM [Colombo et al. 2005] associates metrics with concrete assessment formulas to all the attributes it contains. However, it does not perform well with respect to the first four aspects. Thus, this SQM could be used to annotate QSDs and SLAs which can be used across all service life-cycle activities but in specific scenarios. This is because this SQM contains a rather small amount of domain-independent quality attributes and thus could not be used to capture any possible case in which also some domain-dependent quality attributes are needed. Most of the remaining approaches provide a metric description for some of the included attributes which does not contain a precise assessment formula but an assessment rules set. As the latter rules can be further used to create precise assessment formulas, these approaches could be exploited in all the service life-cycle activities, if extended appropriately.

Based on the previous analysis, no SQM can be considered as optimal according to its evaluation on all the criteria considered. In fact, for particular partitions of the comparison aspects, different approaches are distinguished as the best. Thus, a new SQM is needed that should combine the characteristics of the best approaches in all the considered aspects, describe all the possible but realistic interattribute dependencies, and include also data quality attributes.

6.2. Discussion on Service Quality Metamodels

Many SQMMs have been proposed, which were separated into three partitions based on their scope such that the analysis can be conducted globally for all approaches and locally in each partition. *Pure SQMMs* are able to express QSDs and QSMs. SLA-enabled SQMMs are additionally able to express SLAs. On the other hand, security-based SQMMs focus on particular aspects of service quality description. All SQMMs were evaluated based on a set of criteria capturing the aspects of formality, expressiveness, complexity, and applicability.

Concerning formality, the results have shown that the majority of the approaches use either ontologies or informal formalisms. The former formalism is widely selected in pure SQMMs, while the latter is the best modelling choice in the other two partitions, that is, the SLA-enabled and security-based ones. The analysis has also revealed a recent trend for pure and SLA-enabled SQMMs to use ontologies for their representation. The adoption of ontologies can be explained by their ability to provide unambiguous semantics to quality terms and, thus, to enable machines to automatically process and reason on ontology-specified QSDs in order to support service life-cycle activities like discovery and negotiation.

Three main criteria were used to evaluate the richness of the approaches. The first criterion evaluated the SQM richness of the SQMM. The evaluation results have shown that no SQMM is able to provide a rich SQM. Fortunately, pure SQMMs are starting to improve on this aspect over the last years. However, SLA-enabled and security-based SQMMs do not perform very well on this matter. The second criterion assessed the

richness of the quality metric model and its evaluation results were better with respect to those of the previous one. Indeed, the majority of SQMMs are encompassing an adequately rich quality metric model. Moreover, pure SQMMs are again improving on this matter over the last years. Finally, the third criterion evaluated the richness in constraint description. Here, the evaluation results are even better as the majority of the approaches encompass a rich constraint model. In fact, there are some pure and SLA-enabled SQMMs that encompass an excellent (i.e., very rich) constraint model. Moreover, pure SQMMs are improving on this matter over the years. Security-based SQMMs perform moderately under this evaluation criterion. By closely inspecting the results of these three criteria, it can be inferred that pure SQMMs are continuously increasing their expressiveness, while the approaches of the other two partitions are more or less stable.

It was extremely difficult to assess the SQMM complexity by using good measures based on various practical reasons. Therefore, it was decided to use a simple measure on the number of concepts/entities included in the SQMM and specific thresholds in order to evaluate the SQMMs in particular categories. The results show that most of the SQMMs have low complexity and that there is a trend for pure and SLA-enabled SQMMs to move towards higher complexity. By also considering the fact that SLA-enabled SQMMs are increasing their expressiveness in pure SLA-based aspects, this actually means that modelers are trying to increase the expressiveness of their SQMMs and, in result, the complexity of their SQMMs increases with respect to the number of concepts/entities.

The aspect of applicability was assessed based on two criteria. The first criterion evaluated the connection of an SQMM with a Service Functional Specification Language (SFSL) in order to assess if the SQMM can be used in registries that are bound to specific SFSLS. The results have shown that the majority of the SQMMs are connected to an SFSL. Moreover, the most referenced language is WSDL followed by OWL-S. The second criterion assessed if any service discovery and negotiation framework has adopted the SQMM under inspection. By inspecting the functionality of existing frameworks that use pure SQMMs and considering the fact that these SQMMs do not model some critical information for service monitoring and assessment, it is inferred that pure SQMMs are used until the service negotiation activity. The same result goes for security-based SQMMs. On the other hand, if the SLA-enabled SQMMs are improved on some modelling aspects, then they can be used across the whole service life-cycle.

In summary, from the preceding analysis, there is no SQMM that scores the best value in all criteria. This drawback prevents the wide usage of SQMMs in service management systems. Indeed, as already shown, there are no SQMMs that are used in service discovery, negotiation, and SLA description and enforcement. Thus, there is actually a gap that must be closed by introducing either a new SQMM or extending an appropriate existing one.

6.3. Discussion on Service-Level Agreement Metamodels

Two agreement language types have been proposed in the literature: SLAs and service contracts. The former mainly focus on quality aspects, while the latter have been designed to accommodate for any electronic contract type. Both language types were evaluated on a set of criteria that were grouped along the SLA life-cycle activities. These criteria assessed these languages along the lines of the information they can describe which is required for supporting the SLA life-cycle activities. In this way, by supporting the SLA life-cycle activities, the service life-cycle is also supported.

By inspecting the overall evaluation results, there is no language that satisfactorily supports all activities. On the contrary, for many activities a different language is awarded as the most appropriate one, while only few languages properly support a

subset of all activities. In addition, there are some SLA management activities which are not properly supported by any language, including those of SLA description, match-making, and negotiation.

Concerning SLA languages, the analysis has shown that all SLA languages, including the two most widely used languages, namely WSLA [Keller and Ludwig 2003] and WS-agreement [WS-Agreement 2003], do not possess all appropriate modelling capabilities. Moreover, the capabilities of the WSLA and WS-agreement languages are complementary with respect to the SLA management activities support. For this reason, there are some approaches that attempt to unify the best characteristics of these two languages, such as TrustCom [TrustCoM Consortium 2007]. However, this unification is not enough as those features that are inadequate should be improved and those that are missing should be additionally modelled.

SCLs are not capable of fully supporting most of the SLA management activities apart from those of SLA Monitoring and Assessment and Settlement. This can be explained by the focus of SCL design on service functionality, which was inevitable during SCL modelling time. Thus, although these languages were designed to accommodate for any electronic contract type, they cannot be used to specify SLAs unless they are extended appropriately.

Based on the previous analysis, there is a need for a new language able to express SLAs in a satisfactory way. This language should satisfy all the criteria of all the SLA life-cycle management activities and be capable of explicitly defining SLs, their respective SLOs, and appropriate settlement actions when the SLs are violated or surpassed. The encoding used in this language should enable it to be platform independent, simple, easy to use, and both machine and human understandable and processable. However, the formalism adopted should enable the analysis and the syntactic, semantic, and quality validation of the language's produced SLA specifications, either explicitly or through its transformation to another more powerful formalism. Finally, the high goal of automatability should be achieved with the creation of several assisting tools or SLA management components for this new SLA language that could be used by prospective SPs and SRs, or incorporated in their management systems.

APPENDIX

The Acronyms Used in the Article and their Expansion

ACRONYM	ACRONYM EXPANSION
CG	Constraint Group
CLP	Constraint Logic Programming
FSD	Functional Service Discovery
FSM	Finite State Machine
IaaS	Infrastructure as a Service
QBSM	Quality-Based Service Matchmaking
QoE	Quality of Experience
QoS	Quality of Service
QSD	Quality-Based Service Description
SCL	Service Contract Language
SD	Service Description
SFSL	Service Functional Specification Language
SL	Service Level
SLA	Service Level Agreement
SLA-MM	SLA Meta-Model

ACRONYM	ACRONYM EXPANSION
SLA-SQMM	SLA-Enabled Service Quality Meta-Model
SLO	Service Level Objective
SP	Service Provider
SQM	Service Quality Model
SQMM	Service Quality Meta-Model
SQSL	Service Quality Specification Language
SR	Service Requester
WS	Web Service

REFERENCES

- ALLEN, P. 2006. *Service Orientation, Winning Strategies and Best Practices*. Cambridge University Press, Cambridge, UK.
- ANDRIEUX, A., DAN, A., KEAHY, K., LUDWIG, H., AND ROFRANO, J. 2004. Negotiability constraints in WS-Agreement. Tech. rep. GRAAP-WG. Version 0.1. <http://www.mcs.anl.gov/~keahey/Meetings/GRAAP/WSAgreement%20Negotiability%20Constraints.pdf>.
- BATINI, C., CAPIELLO, C., FRANCALANCI, C., AND MAURINO, A. 2009. Methodologies for data quality assessment and improvement. *ACM Comput. Surv.* 41, 3.
- BRANDIC, I., BUYYA, R., MATTESS, M., AND VENUGOPAL, S. 2009. Towards a meta-negotiation architecture for sla-aware grid services. In *Proceedings of the 2nd International Workshop on Service Oriented Engineering and Optimization (SENOPT'08) in conjunction with Conference on High Performance Computing*. 1–9.
- BRANDIC, I., PLLANA, S., AND BENKNER, S. 2006. An approach for the high-level specification of qos-aware grid workows considering location affinity. *Sci. Programm. J.* 14, 3–4, 231–250.
- CAPPIELLO, C. 2006. The quality registry. In *Mobile Information Systems – Infrastructure and Design for Adaptivity and Flexibility*, Springer, 307–317.
- CAPPIELLO, C., KRITIKOS, K., METZGER, A., PARKIN, M., PERNICI, B., PLEBANI, P., AND TREIBER, M. 2008. A quality model for service monitoring and adaptation. In *Proceedings of the Workshop on Monitoring, Adaptation and Beyond (MONA+) at the ServiceWave Conference*. Springer.
- COLOMBO, M., NITTO, E. D., PENTA, M. D., DISTANTE, D., AND ZUCCAL, A. M. 2005. Speaking a common language: A conceptual model for describing service-oriented systems. In *Proceedings of the 3rd International Conference on Service Oriented Computing (ICSOC'05)*. 48–60.
- COMUZZI, M., KRITIKOS, K., AND PLEBANI, P. 2009. A semantic based framework for supporting negotiation in service oriented architectures. In *Proceedings of the 11th IEEE Conference on Commerce and Enterprise Computing (CEC'09)*. 137–145.
- CORTÉS, A. R., MARTÍN-DÍAZ, O., TORO, A. D., AND TORO, M. 2005. Improving the automatic procurement of web services using constraint programming. *Int. J. Cooperative Inf. Syst.* 14, 4, 439–468.
- CRANOR, L., DOBBS, B., EGELMAN, S., HOGBEN, G., HUMPHREY, J., LANGHEINRICH, M., MARCHIORI, M., PRESLERMARSHALL, M., REAGLE, J., SCHUNTER, M., STAMPLEY, D. A., AND WENNING, R. 2006. Platform for privacy preferences (p3p). Working group note, W3C. November. <http://www.w3.org/P3P/>.
- DE PAOLI, F., PALMONARI, M., COMERIO, M., AND MAURINO, A. 2008. A meta-model for non-functional property descriptions of web services. In *Proceedings of the IEEE International Conference on Web Services (ICWS'08)*. 393–400.
- DIKAIAKOS, M. D., PALLIS, G., KATSAROS, D., MEHRA, P., AND VAKALI, A. 2009. Cloud computing: Distributed internet computing for it and scientific research. *IEEE Internet Comput.* 13, 5, 10–13.
- DOBSON, G., LOCK, R., AND SOMMERVILLE, I. 2005. QoSOnt: A qos ontology for service-centric systems. In *Proceedings of the 31st EUROMICRO Conference on Software Engineering and Advanced Applications (EUROMICRO'05)*. 80–87.
- FARRELL, A. D. H., SERGOT, M. J., TRASTOUR, D., AND CHRISTODOULOU, A. 2004. Performance monitoring of service-level agreements for utility computing using the event calculus. In *Proceedings of the 1st IEEE International Workshop on Electronic Contracting (WEC'04)*. 17–24.
- FOSBROOK, D. AND LAING, A. C. 1996. *The A-Z of Contract Clauses*. Sweet and Maxwell.
- FROLUND, S. AND KOISTINEN, J. 1998. Quality of services specification in distributed object systems design. In *Proceedings of the 4th Conference on USENIX Conference on Object-Oriented Technologies and Systems (COOTS'98)*. 179–202.

- FRUTOS, H. M., KOTSIPOULOS, I., GONZALEZ, L. M. V., AND MERINO, L. R. 2009. Enhancing service selection by semantic qos. In *Proceedings of the 6th European Semantic Web Conference on The Semantic Web: Research and Applications (ESWC'09)*. 565–577.
- GEORGAKOPOULOS, D. AND PAPAIOGLOU, M. P., EDS. 2008. *Service-Oriented Computing*. MIT Press. http://mitpress.mit.edu/sites/default/files/titles/content/9780262072960_sch.0001.pdf.
- GIALONARDO, E. AND ZIMEO, E. 2007. More semantics in qos matching. In *Proceedings of the International Conference on Service-Oriented Computing and Applications*. 163–171.
- GREFEN, P. AND ANGELOV, S. 2002. On τ -, μ -, π -, and ε -contracting. In *Proceedings of the International Workshop on Web Services, E-Business, and the Semantic Web (WES'02)*. Lecture Notes in Computer Science, vol. 2512, Springer, 68–77.
- GROSO, B. N. AND POON, T. C. 2004. SweetDeal: Representing agent contracts with exceptions using semantic web rules, ontologies, and process descriptions. *Int. J. Electron. Commerce* 8, 4, 61–97.
- HOFFNER, Y., FIELD, S., GREFEN, P., AND LUDWIG, H. 2001. Contract-driven creation and operation of virtual enterprises. *Comput. Netw.* 37, 111–136.
- HWANG, C. AND YOON, K. 1981. Multiple criteria decision making. Lecture Notes in Economics and Mathematical Systems. Tech. rep. ISO/IEC 2001. ISO/IEC 9126-1 Software Engineering. Product Quality - Part 1: Quality model. ISO/IEC.
- JIANG, Y., SHAO, W., ZHANG, L., MA, Z., MENG, X., AND MA, H. 2004. On the classification of umls meta model extension mechanism. In *Proceedings of the 7th International Conference on The Unified Modeling Language: Modeling Languages and Applications (UML'04)*. 54–68.
- KAZHAMIAKIN, R., PISTORE, M., AND ZENGIN, A. 2009. Cross-layer adaptation and monitoring of service-based applications. In *Proceedings of the International Conference on Service-Oriented Computing (ICSOC/ServiceWave'09)*. Lecture Notes in Computer Science, vol. 6275, Springer, 325–334.
- KELLER, A. AND LUDWIG, H. 2003. The wsla framework: Specifying and monitoring service level agreements for web services. *J. Netw. Syst. Manag.* 11, 1, 57–81.
- KRITIKOS, K. 2008. QoS-based web service description and discovery. PhD thesis, Computer Science Department, University of Crete, Heraklion, Greece. December. <http://ercim-news.ercim.eu/qos-based-web-service-description-and-discovery>.
- KRITIKOS, K. AND PLEXOUSAKIS, D. 2006. Semantic qos metric matching. In *Proceedings of the European Conference on Web Services (ECOWS'06)*. 265–274.
- KRITIKOS, K. AND PLEXOUSAKIS, D. 2009. Requirements for qos-based web service description and discovery. *IEEE Trans. Serv. Comput.* 2, 4, 320–337.
- LAMANNA, D. D., SKENE, J., AND EMMERICH, W. 2003. SLAng: A language for defining service level agreements. In *Proceedings of the 9th IEEE International Workshop on Future Trends of Distributed Computing Systems (FTDCS'03)*.
- LININGTON, P. F., MILOSEVIC, Z., COLE, J., GIBSON, S., KULKARNI, S., AND NEAL, S. 2004. A unified behavioural model and a contract language for extended enterprise. *Data Knowl. Engin.* 51, 1, 5–29.
- MA, Y., JIN, B., AND FENG, U. 2010. Semantic-oriented ontology cohesion metrics for ontology-based systems. *J. Syst. Softw.* 83, 1, 143–152.
- MABROUK, N. B., GEORGANTAS, N., AND ISSARNY, V. 2009. A semantic end-to-end qos model for dynamic service oriented environments. In *Proceedings of the Principles of Engineering Service Oriented Systems collocated with the International Conference on Software Engineering*.
- MAXIMILIEN, E. M. AND SINGH, M. P. 2002. Conceptual model of web service reputation. *SIGMOD Rec.* 31, 4, 36–41.
- MAXIMILIEN, E. M. AND SINGH, M. P. 2004. A framework and ontology for dynamic web services selection. *IEEE Internet Comput.* 8, 5, 84–93.
- MENS, T. AND LANZA, M. 2002. A graph-based metamodel for object-oriented software metrics. *Electron. Notes. Theor. Comput. Sci.* 72, 2.
- MOLINA-JIMENEZ, C., SHRIVASTAVA, S., SOLAIMAN, E., AND WARNE, J. 2003. Contract representation for run-time monitoring and enforcement. In *Proceedings of the IEEE International Conference on E-Commerce Technology (CEC'03)*. 103–110.
- MÜLLER, C., CORTÉS, A. R., AND RESINAS, M. 2008. An initial approach to explaining sla inconsistencies. In *Proceedings of the 6th International Conference on Service-Oriented Computing (ICSOC'08)*. Lecture Notes in Computer Science, vol. 5364, Springer, 394–406.
- MULLER, N. J. 1999. Managing service level agreements. *Int. J. Netw. Manag.* 9, 3, 155–166.
- NADALIN, A., GOODNER, M., GUDGIN, M., BARBIR, A., AND GRANQVIST, H. 2007. WS-trust specification, Tech. rep. OASIS working draft. <http://www.ibm.com/developerworks/webservices/library/specification/ws-trust/>.

- NEJDL, W., OLMEDILLA, D., AND WINSLETT, M. 2004. PeerTrust: Automated trust negotiation for peers on the semantic web. In *Proceedings of the VLDB International Workshop on Secure Data Management in a Connected World (SDM'04)*. Lecture Notes in Computer Science, vol. 3178, Springer, 118–132.
- NESSI OPEN FRAMEWORK. 2009. Quality model for nexof-ra pattern designing. Tech. rep. http://ec.europa.eu/information_society/apps/projects/logos/6/216446/080/deliverables/001.D10220100901.pdf.
- THE OASIS GROUP. 2005. Quality model for web services. Tech. rep., The Oasis Group. <http://www.clip.dia.fi.upm.es/Projects/S-CUBE/papers/oasis05:WSQM-2.0.pdf>.
- THE OMG GROUP. 2005. UMLTM profile for modeling quality of service and fault tolerance characteristics and mechanisms. Tech. rep. ptc/2005-05-02, The OMG Group.
- OLDHAM, N., VERMA, K., SHETH, A., AND HAKIMPOUR, F. 2006. Semantic ws-agreement partner selection. In *Proceedings of the 15th International Conference on World Wide Web (WWW'06)*. ACM Press, New York, 697–706.
- OREN, N., PREECE, A., AND NORMAN, T. 2005. Service level agreements for semantic web agents. In *AAAI Fall Symposium Series*.
- O'SULLIVAN, J., EDMOND, D., AND TER HOFSTEDÉ, A. 2002. What's in a service? Towards accurate description of non-functional service properties. *Distrib. Paralle. Datab.* 12, 2–3, 117–133.
- PARKIN, M., BADIA, R. M., AND MARTRAT, J. 2008. A comparison of sla use in six of the european commissions fp6 projects. Tech. rep. TR-0129, Institute on Resource Management and Scheduling, CoreGRID - Network of Excellence.
- PASCHKE, A. 2005. RBSLA: A declarative rule-based service level agreement language based on ruleml. In *Proceedings of the International Conference on Computational Intelligence for Modelling, Control and Automation and International Conference on Intelligent Agents, Web Technologies and Internet Commerce (CIMCAIAWTIC'06)*. Vol. 2, 308–314.
- PASCHKE, A. AND SCHNAPPINGER-GERULL, E. 2006. A categorization scheme for sla metrics. In *Service Oriented Electronic Commerce: Proceedings zur Konferenz im Rahmen der Multikonferenz Wirtschaftsinformatik*. LNI, vol. 80, 25–40.
- RAN, S. 2003. A model for web services discovery with QoS. *SIGecom Exch.* 4, 1, 1–10.
- REDMAN, T. C. 1997. *Data Quality for the Information Age*. Artech House, Inc., Norwood, MA.
- ROSSI, F., VAN BEEK, P., AND WALSH, T. 2006. *Handbook of Constraint Programming (Foundations of Artificial Intelligence)*. Elsevier Science, New York.
- SABATA, B., CHATTERJEE, S., DAVIS, M., SYDIR, J., AND LAWRENCE, T. 1997. Taxonomy for qos specifications. In *Proceedings of the 3rd International Workshop on Object-Oriented Real-Time Dependable Systems*. 100–107.
- SAKELLARIOU, R. AND YARMOLENKO, V. 2008. Job scheduling on the grid: Towards sla-based scheduling. <http://www.cs.man.ac.uk/~rizos/papers/hpc08.pdf>.
- SETH, N., DESMUKH, S. G., AND VRAT, P. 2005. Service quality models: A review. *Int. J. Qual. Reliab. Manag.* 22, 9, 913–949.
- SKENE, J. 2007. Language support for service-level agreements for application-service provision. Ph.D thesis, Department of Computer Science, University College London, London, UK. <http://discovery.ucl.ac.uk/5607/>.
- SKOGSRUD, H., BENATALLAH, B., AND CASATI, F. 2004. Trust-serv: Model-driven lifecycle management of trust negotiation policies for web services. In *Proceedings of the 13th Conference on World Wide Web*.
- STRONG, D. M., LEE, Y. W., AND WANG, R. Y. 1997. 10 pitholes in the road to information quality. *Comput.* 30, 8, 38–46.
- TEBBANI, B. AND AIB, I. 2006. GXLA a language for the specification of service level agreements. In *Proceedings of the 1st International IFIP TC6 Conference on Autonomic Networking*. Lecture Notes in Computer Science, vol. 4195, Springer, 201–214.
- TIAN, M., GRAMM, A., NABULSI, M., RITTER, H., SCHILLER, J., AND VOIGT, T. 2003. QoS integration in web services. Gesellschaft für Informatik DWS 2003, Doktorandenworkshop Technologien und Anwendungen von XML.
- TOSIC, V., ESFANDIARI, B., PAGUREK, B., AND PATEL, K. 2002. On requirements for ontologies in management of web services. In *Revised Papers from the International Workshop on Web Services, E-Business, and the Semantic Web (CAiSE'02/WES'02)*. Springer, 237–247.
- TOSIC, V., MA, W., PAGUREK, B., AND ESFANDIARI, B. 2003a. On the dynamic manipulation of classes of service for xml web services. Research rep. SCE-03-15, Department of Systems and Computer Engineering, Carleton University, Ottawa, Canada.
- TOSIC, V. AND PAGUREK, B. 2005. On comprehensive contractual descriptions of web services. In *Proceedings of the IEEE International Conference on e-Technology, e-Commerce and e-Service (EEE'05)*. 444–449.

- TOSIC, V., PAGUREK, B., AND PATEL, K. 2003b. WSOL – A language for the formal specification of classes of service for web services. In *Proceedings of the International Conference on Web Services (ICWS'03)*, L.-J. Zhang, Ed. CSREA Press, Las Vegas, NV, 375–381.
- TRUONG, H.-L., SAMBORSKI, R., AND FAHRINGER, T. 2006. Towards a framework for monitoring and analyzing qos metrics of grid services. In *Proceedings of the International Conference on e-Science and Grid Computing*.
- TRUSTCOM CONSORTIUM. 2007. TrustCom framework v4 – Appendix a: Profiles. Report Deliverable D63, European Union.
- WANG, X., VITVAR, T., KERRIGAN, M., AND TOMA, I. 2006. A qos-aware selection model for semantic web services. In *Proceedings of the 4th International Conference on Service-Oriented Computing (ICSOC'06)*, A. Dan and W. Lamersdorf, Eds. Lecture Notes in Computer Science, vol. 4294, Springer, 390–401.
- WELTY, C., KALRA, R., AND CHU-CARROLL, J. 2003. Evaluating ontological analysis. In *Proceedings of the ISWC-03 Workshop on Semantic Integration*.
- WS-AGREEMENT. 2003. WS-agreement framework. <https://forge.gridforum.org/projects/graap-wg>.
- YI, T., WU, F., AND GAN, C. 2004. A comparison of metrics for uml class diagrams. *SIGSOFT Softw. Engin. Notes* 29, 5, 1–6.
- YOA, H., OREM, A. M., AND ETZKORN, L. 2005. Cohesion metrics for ontology design and application. *J. Comput. Sci.* 1, 1, 107–113.
- ZEMNI, M. A., BENBERNOU, S., AND CARRO, M. 2010. A soft constraint-based approach to qos-aware service selection. In *Proceedings of the 8th International Conference on Service Oriented Computing (ICSOC'10)*. Springer.
- ZHOU, C., CHIA, L.-T., AND LEE, B.-S. 2004. DAML-qos ontology for web services. In *Proceedings of the IEEE International Conference on Web Services (ICWS'04)*. 472–479.

Received November 2009; revised July 2012; accepted October 2012