



What's in a Service?*

Towards Accurate Description of Non-Functional Service Properties

JUSTIN O'SULLIVAN

justin.osullivan@gbst.com

GBST Holdings Pty Ltd, P.O. Box 1511, Milton QLD 4064, Australia; Centre for Information Technology Innovation, Queensland University of Technology, GPO Box 2434, Brisbane QLD 4001, Australia

DAVID EDMOND AND ARTHUR TER HOFSTEDE

Centre for Information Technology Innovation, Queensland University of Technology, GPO Box 2434, Brisbane QLD 4001, Australia

Recommended by: Boualem Benatallah and Fabio Casati

Abstract. A proper understanding of the general nature, potential and obligations of electronic services may be achieved by examining existing commercial services in detail. The everyday services that surround us, and the ways in which we engage with them, are the result of social and economic interaction that has taken place over a long period of time. If we attempt to provide electronic services, and do not take this history into account, then we will fail. Any attempt to provide automated electronic services that ignores this history will deny consumers the opportunity to negotiate and refine, over a large range of issues, the specific details of the actual service to be provided. To succeed, we require a rich and accurate means of representing services. An essential ingredient of service representation is capturing the non-functional properties of services. These include the methods of charging and payment, the channels by which the service is requested and provided, constraints on temporal and spatial availability, service quality, security, trust and the rights attached to a service. Not only are comprehensive descriptions essential for useful service discovery, they are also integral to service management, enabling service negotiation, composition, and substitution. This paper builds on an understanding of services and their interactions, to outline the non-functional properties of services and their uses.

Keywords: service description, service properties, service substitution, electronic services

1. Introduction

Services are ubiquitous, yet there is no adequate standard for accurately describing them. The need to describe a service is like the need to label goods or products in a supermarket. A product label provides a brief summary of the good to which it is attached. Prospective buyers can use this information, together with the price, to make a rational purchasing decision. Product labelling occurs for the safety and benefit of purchasers. Why is the same accurate labelling (or description) not provided for the benefit of service consumers?

*This work is funded by an Australian Research Council SPIRT Grant entitled "Self-describing transactions operating in a large, open, heterogeneous and distributed environment" involving QUT and GBST Holdings Pty Ltd.

Inadequate description triggers questions for a requestor. When you encounter a service how do you request it? Where and when is the service available? By what means do you access the service? What quality of service are you guaranteed? What is the identity of the service provider? What payment and settlement models are available? What rights do you have over the service? Where is the manifest that describes how the service is composed? Once requested, what are the models of interaction that might occur during its delivery? Current initiatives such as Universal Description, Discovery and Integration (UDDI), the Web Services Description Language (WSDL) and the Web Services Flow Language (WSFL) are attempting to provide a basis for business integration [12, 18, 30]. We argue that these initiatives lack the accuracy required to operate on services. Accurate service description will benefit the following activities:

- *Discovery*: more rapid and accurate service discovery will be possible than is permitted by existing catalogues or registries.
- *Substitution*: accurate service descriptions form a basis for comparing services, enabling rational optimisation and negotiation, for example.
- *Composition*: service descriptions may be used to build new services, either statically or dynamically.
- *Management*: based upon the rich service description languages and operations, architectures may be developed for managing repositories of service metadata. This metadata could be used for many purposes including controlling and monitoring the discovery, substitution and composition processes.

Whilst we acknowledge the importance of service functionality, this paper is primarily concerned with the non-functional properties of services. A service is *not* a function. It is a function performed on your behalf at a cost. And the cost is not just some monetary price; it is a whole collection of limitations. This paper is all about these. We consider the non-functional properties to be constraints over the functionality of the service [13]. We believe that a service description is only complete once the non-functional aspects are also expressed. Service description initiatives now need to be focused on determining the particular properties that should be represented and the degree of representation that is required. This paper primarily focuses on the former of these issues. The sections that follow attempt to define services, present their associated interactions, and more importantly detail the non-functional properties (i.e. temporal and spatial availability, channels, charging styles, settlement models and contracts, payment, quality, security, trust and rights) that comprise services. Finally we offer an overview of how the non-functional properties of services can be used to assist the service life cycle before presenting our conclusions.

2. What is a service?

Many definitions for services are based on technology. Some definitions of *electronic* services (or e-services) use the Internet and/or workflow as a conduit to new revenue or task completion [40, 42, 50]. A *web* service has been described as an aggregation of functionality published for use [27]. This is similar to the virtual business processes that define

company-level interactions [28]. Other definitions offer a view of services as an abstraction of a business process [10, 25, 45]. We assert that e-services exhibit minimal constraints on the time and location of *request*. There may, however, be a delay between the request and the *execution*. Such a delay may arise because of resource constraints or because of human intervention required in the performance of the service. We believe that an e-service is characterised by its ability to be automatically *summoned* anywhere, anytime.

Surrounding these definitions are three immutable features of services. Firstly, that services are actions performed by an entity on behalf of another. Secondly, services are an asset [40]. They have an inherent value that is transferred from the provider to the recipient. Finally, services can be contained within other services (e.g. a warranty) [53]. We refer to this relationship by describing some services as *sub-services*.

We consider services that contain other services (or sub-services) to be either an aggregation or a composition. Aggregations combine multiple services and provide access to them in a single location. Telecommunications companies can be considered an example of service aggregators. Services such as call forwarding, call diversion and voicemail, are brought together and offered via the telephone. A composition is a tightly-coupled *integration* of sub-services that results in value not present within the individual services. This added value may be represented in terms of another service property (e.g. reduced price, increased trust). Within a composition, each sub-service is a service in its own right and complex inter-relationships may exist between the sub-services. Service composition should not be confused with functional composition. It has a broader goal that needs to take into account both functional *and* non-functional issues. It may be that we can (functionally) compose a transportation service by articulating land and air transport services. We may equally (non-functionally) compose some hitherto free service with a payment mechanism to form a commercial version of the original service.

3. Service interactions

Service interactions include three primary participants, a service provider, a service requestor and a service catalogue (or registry) [10, 11, 17, 27]. A more detailed view has been offered in [21]. A fourth participant, the service broker, is sometimes present in service interactions. We consider service brokers (e.g. an insurance broker) to be entities that offer services from multiple providers to a requestor. Service brokers attempt to add value to the service provision process, sometimes protecting the identity of the service providers for their own benefit. We identify four main interactions between these participants: discovery, negotiation, invocation, and execution.

The recognition of a need, by a service requestor, triggers the search for suitable service providers. Requestors with little or no knowledge of appropriate providers normally interact with service catalogues. Catalogues are themselves normally well-known services that compare the needs of a requestor against the advertised service descriptions it maintains. The names of candidate providers, if available, are returned to the requestor. Requestors who are aware of the appropriate provider(s) may bypass the provider search with a catalogue and directly approach the provider. To be able to *discover* a service requires that it be defined, somehow, and that this definition be published.

Requestors contact candidate providers and undertake a series of interactions that refine their knowledge about the functionality and the non-functional properties (e.g. availability and quality) of each potential service. Generally, this *negotiation* results in a “service contract” that outlines the obligations of each party, and that may or may not be enacted. Requestors have the opportunity to refine their knowledge with respect to multiple candidate providers before making the decision to enact a particular service.

Invocation is the term we use to identify the process which begins with the formation of a binding agreement between requestor and provider. This is essentially a *call* for the execution of a service. Invocation of the service contract also triggers the provisioning and/or production of the service by the provider. Services can be invoked using different forms of binding requests. These include electronic requests (e.g. URI or web page), verbal requests (e.g. buying from a shop), written requests (e.g. invoice or letter), manual actions (e.g. turning on the TV) and sensor-based requests (e.g. automated door, barcode scanner).

We refer to the delivery and consumption of a service as its *execution*, which normally results in the fulfilment of each party’s obligations. Some services are delivered at a location and time distinct from the invocation. For example, when booking a trip you might walk to or phone your travel agent. The trip is booked and the service will be rendered at the airport, at the nominated date and time. A specific instance of service delivery can occur within the context of an existing service. This is evident when you catch a bus that is moving between points A and Z. You request the service (i.e. hail and get on the bus) at point D and your consumption of the service ends at point G (i.e. you get off). In the example provided it is interesting to note that the service provider (i.e. the bus company) may deliver the service and it is never consumed (i.e. the bus is driven from A to Z and nobody gets on). This is the characteristic of perishability described in [53].

Consumption may involve the suspension and resumption of the service. Consumption is also an optional aspect of a service (e.g. you receive a monthly newsletter and by throwing it directly in the bin you are not consuming it).

4. Non-functional properties

We now present a discussion of the *non-functional* properties associated with services. As previously mentioned we consider these non-functional properties to be constraints exhibited over the functionality of the service. The non-functional properties of services include temporal and spatial availability, channels, charging styles, settlement models, settlement contracts, payment, service quality, security, trust and ownership. Each of the properties presented in this section is deserving of a separate paper. Instead, we try to reveal sufficient detail with respect to each non-functional property as to outline the complexity involved with accurately describing them. The non-functional properties outlined in this paper have been based on a review of existing commercial services. Other properties may be discovered during the course of our research.

4.1. Availability

We consider *availability* to refer to the temporal (i.e. when) and spatial (i.e. where) constraints applied to a service. Availability is a complex property of services. For example,

there are services that are regularly on the move (e.g. taxis, trains). There are also services where an implicit understanding affects the advertised availability (e.g. when attending the theatre, you need to be in the lobby prior to the start time so that seating can take place.) Thirdly, there are services where there is a suspension and resumption (e.g. memberships).

4.1.1. Temporal and spatial issues. In representing complex spatio-temporal information other issues, apart from those outlined, need to be addressed. Often, services quite intentionally provide incomplete spatio-temporal information. For example, when you buy an airline ticket you know the airport where the plane departs from. Further refinement occurs at check-in to include a departure gate number, boarding time and a seat number. Availability of a service may be specified with respect to another object (e.g. an emergency phone is available 3 km south of a particular overpass on the freeway). This is also referred to as orientation and is defined using the primary object, a reference object and a frame of reference [15]. Different temporal representations can be used but they assume a “degree of certainty” about the information being represented [3]. Uncertainty increases with a reduced frequency of sampling [39, 43]. How do we know that a bus will arrive according to its timetable? We can assume that it will arrive on time or we can stand at the bus stop and continually check. It is important that uncertainty is communicated to the service requestor. Some services have exclusivity arrangements relating to their availability (e.g. an appointment for the doctor or hire of a conference centre). Location-based services (e.g. where is the nearest hotel to where I am now?) also face the same representational challenges.

For decision-making reasons, service requestors may need to be aware of more than just the availability of service request and delivery times. To enable accurate scheduling of multiple services, the requestor may be specifically interested in the duration of the service or the approximate completion time. These may be required when performing service discovery, advertising, composition, and when determining service quality.

4.1.2. Temporal and spatial representation. Temporal representations need to support various granularities or alternatively represent time as a relationship (e.g. service “X” begins after service “Y”). Common temporal granularities include seconds, minutes, hours, days, weeks, months and years. Approaches for capturing these granularities and their relationships (e.g. finer-than, groups-into) have previously been offered [1, 7]. Temporal database literature has well-defined terms such as chronon (non-decomposed unit of time), time stamp, lifespan, event and interval [26]. Analogous to chronons is the concept of a moment [2]. Each of these concepts offers insight into the expression of granularities for temporal availability. A useful summary of the problems associated with using temporal time stamps such as *now* are outlined in [14]. Another method for representing date and time is the ISO standard 8601:2000 [24], which is intended for use in software to software exchanges.

The artificial intelligence community uses dating schemes, constraint propagation and duration-based schemes for temporal representation [3]. Within the spatio-temporal database community sets of object, location, and time-stamp triplets have been used to represent time evolving spatial objects [48]. Three temporal specification issues are outlined, each of which is applicable to services: (1) data type support for service definition languages; (2) index construction for service catalogs; and (3) query processing for service discovery.

Spatial representations are used to describe topologies, orientation, shape, size and distance [15]. A discussion of spatial models and their classifications (comprehensiveness, structure, theoretical foundation, modelling techniques) is found in [22]. Latitude, longitude and altitude (e.g. for planes) may also be useful for describing services. Representation and indexing of moving-point objects is discussed in detail in [39, 43]. Service routes (e.g. a bus route), and service regions (e.g. airports) will require spatial representation.

Filtering is sometimes applied to limit the spatial availability of a service to some requestors. Some examples of filtering include calling a phone number that redirects the requestor to the appropriate provider in your region, or franchises that operate only within a specific suburb(s).

4.2. Channels

We consider service interactions to occur using a channel [20]. When attempting to describe a channel we need to take into consideration its endpoints, the information being transmitted and the interaction pattern that occurs over the channel. We categorise service description initiatives into those that describe service functionality or non-functional properties, and those that describe interactions with services.

Functionality and property description initiatives include the WSDL, the DARPA Agent Markup Language for Web Services (DAML-S), IBM's Web Service Endpoint Language (WSEL), Corba's Component Definition Language (CDL), and IBM's Advertisement and Discovery of Services (ADS) protocol [4, 12, 30, 35, 51]. Of these, WSDL and DAML-S are the most relevant. WSDL attempts to describe channel endpoints (i.e. their interface, protocol, bindings and operations). DAML-S provides a web service ontology aimed at increasing the automation of services. Our work compliments DAML-S by providing a foundation for the non-functional properties that require representation. All of the above mentioned initiatives fail to fully address non-functional properties such as quality, settlement models, settlement contracts, trust and rights.

Interaction initiatives are currently focused on representing choreographed dialogues between service providers and service requestors. Generic languages include WSFL, the Web Services Conversation Language (WSCL) and XLANG. IBM's WSFL is a workflow-like approach to assembling services [30]. WSCL is HP's attempt to describe the legal sequence of messages that can be exchanged between entities [6]. XLANG which is built on WSDL, is used to describe business processes [47]. Numerous business-level initiatives exist within this area (e.g. RosettaNet Partner Interface Processes [19] and BizTalk [38]) and we consider a comprehensive survey of these frameworks outside the scope of this paper. We consider the ability to specify a service dialogue as a necessary component for these initiatives.

Interactions may occur over channels utilising broadcast techniques. This is a means of addressing an unknown number of providers (e.g. placing a wanted advertisement in the *classifieds* section of a newspaper) or requestors (e.g. receiving news updates from a web site). This technique is more common with delivery channels and is commonly referred to as *pushing*. Broadcast channels have the unique property that they may not have been

explicitly requested (e.g. a television or radio station) and additionally, may have no associated request channel.

Services accessible by two different channels may offer overlapping or distinct operations (e.g. a bank balance can be requested over the counter or via the Internet, however a bank cheque can only be drawn over the counter). Channels may support varying properties or characteristics such as temporal availability or price. Separation of the request and delivery channels may introduce the need to acknowledge a request. In some cases this is provided in the form of a receipt.

4.3. *Charging styles*

The styles presented here describe the charging technique applied by a service provider for the use of its service. Three styles are identified: (1) per service request or delivery (e.g. a fixed price local telephone call); (2) by unit of measure and granularity (e.g. by length, volume, weight, area or time); and (3) on a percentage or ratio basis of some aspect of the service (e.g. by commission).

Service providers may use an aggregation of charging styles. An example of this is a telecommunications provider (e.g. AT&T, Deutsche Telecom). The services of a telco are charged using multiple styles. This includes granular services such as per minute or second phone calls (either interstate, international or mobile phone) and per month line rental. Charges such as the initial connection fee and fixed cost local phone calls are charged on a per service request basis.

Sometimes the charge for a service is redirected to another entity. An example of this is a free web-based email service. No cost is applied to the service requestor but advertising is used to pay for the service.

4.4. *Settlement*

Settlement is a process that reflects the mutual obligations of the provider and requestor, with payment usually being an obligation of the requestor, and service delivery being that of the provider. The settlement process is normally laid down by the provider, and is included as part of their business model. The settlement process (and its sequence) is sometimes defined by the service environment.

4.4.1. *Settlement models.* Packaging of obligations into a defined process results in a settlement model, which reflects the ordering and relationship between each party's obligations. None of the settlement models presented in this section result in a transfer of ownership (see Section 4.8). The service requestor does however consider some form of value to be transferred during the provision process.

Two well-known settlement models are the *transactional* and the *rental* models. The transactional model can be described simply as delivery for payment. It can be a one-off delivery or include multiple deliveries of the same service. The later implies a longer term relationship. The rental model is the familiar concept of being "on loan" (e.g. a video). Within the rental model, explicit temporal or spatial constraints may be imposed by the

service provider (e.g. (a) the video is to be returned by 6 pm tomorrow, or (b) when hiring a conference centre the service is found at a physical address). Depending on the service, rental may involve a short-term relationship (e.g. holiday unit) or long-term relationship (e.g. local video store membership).

Specialised forms of the transactional model are (1) *subscription*, which normally implies a long-term relationship; (2) *metered*, which is almost identical to the basic transactional model, tracks consumption of the service except that the relationship may also impose restrictions making it difficult to change to another service provider; (3) *facilitated*, in which the provider acting as a conduit or facilitator to another service provider (e.g. broker or financial planner); (4) *escrow*, which is used when there is an identified trust issue, and where the parties lodge their obligation with the escrow organisation; and (5) *swap*, where the parties agree that the services being traded are of equal value, and no payment is involved.

Multi-party settlements vary in their degree of binding between the parties involved. A tightly bound third-party might include a credit card provider (e.g. a bank) and a loosely bound third-party might include a company that provides software used during service provision (e.g. an accounting package).

Service providers sometimes trigger the obligations of the service requestor by using a request for payment or an invoice. This may indicate that the service provider has completed its obligations.

4.4.2. Settlement contracts. Attached to the service may be a list of terms and conditions. These conditions are formalised in a contract and govern the responsibilities of all parties involved in the service request and provision. Contracts are considered binding agreements between parties [9]. Types of settlement contracts in an offline environment include (1) a *Bill of Lading*, which defines details of transportation (e.g. who, what, where) and what happens should something go wrong; and (2) a *Promissory Note*, which outlines the terms and conditions of a loan (e.g. required repayments, interest rate and policies surrounding the loan). More familiar examples include the terms and conditions associated or expressed with items such as credit card applications, tickets for transportation or entertainment, and policies (e.g. insurance). We assert that these contracts are representations of the promises of each party. Both parties must be agreeable with respect to the contract before it is invoked. An implementation of the infrastructure required for electronic contracts is outlined in [23]. In the case of a warranty, contracts may prove to extend the life of a service beyond the initial transfer of value.

Recourse is available in some cases to either the service provider or the service requestor. In cases where obligations of either party are not realised there may be some level of re-negotiation performed. A contracting protocol that includes the ability to decommit is outlined in [44].

4.5. *Payment obligations*

Payment obligations may be required at any stage (e.g. upfront, in arrears, staged installments) in the service invocation, provision and execution process. These obligations are

normally outlined to the service requestor as part of the negotiation process and are included in any attached settlement contracts. Service providers or their surrounding environment determine a valid set of payment instruments that are used to fulfil this important obligation of the service requestor. Payment instruments are used within the context of a payment model. The entities and information flows associated with payment models have previously been outlined in [37]. Additionally we recognise that payment protocols (e.g. Society for Worldwide Interbank Financial Telecommunications [46]) are sometimes used as a mechanism for controlling the flows within these models.

We consider the term payment instruments to be relatively self-explanatory. We consider payment instruments to include items such as cash, cheques, direct funds transfers, credit or charge cards, travellers' cheques, wire transfers, postal or money orders, securities (i.e. stocks, options, warrants), bank bills, vouchers, stored value cards, digital cash and anonymous cash. A useful summary of payment instrument dimensions is provided in [31].

4.6. *Service quality*

Service quality is a measure of the difference between expected and actual service provision. It is a complex and largely domain-specific property. From the viewpoint of the requestor, it measures the competence of the provider to deliver a service [29, 41]. The most notable work on measuring customer perceptions of service quality is SERVQUAL [36]. This work produced scale that measured perceived service quality along five dimensions: the dependability and accuracy of the service (*reliability*); the promptness and the willingness of staff to assist (*responsiveness*); attributes, such as knowledge and courtesy, of staff that conveyed trust and confidence to the user (*assurance*); the level of caring and personalised attention provided to the requestor (*empathy*); and concrete or physical aspects of the service, such as cleanliness (*tangibles*).

Service providers may commit to providing a certain level of quality. This commitment is sometimes formalised using a *Service Level Agreement*. Service level agreements can be considered as binding contracts that are agreed between a service provider and service requestor. Penalties are normally imposed for non-compliance. Commitment to a service can be bound into the contracting protocol [44]. This offers a method of backing out of a service, assuming that the agreed penalty is paid. Service providers also use guarantees or warranties to express commitment to a service. A useful survey of service quality frameworks is outlined in [5].

4.7. *Security and trust*

Security and trust are foundational properties for electronic service provision. The issues surrounding their use in services are outlined below. We do not discuss possible representations for either of these topics.

4.7.1. Security. Security is increasingly being viewed as a mandatory component for facilitating electronic commerce. It alleviates concerns relating to identity, privacy, alteration and repudiation of information transferred between parties [8]. We commonly think about

“on-the-wire” security that pertains to the request and delivery channels of a service, especially when the payment obligation of the service requestor is being finalised. Security protocols such as the Secure Sockets Layer are becoming widespread for this role. Common approaches to security within organisations involve the implementation of a Public Key Infrastructure (PKI).

We believe that individual aspects of service descriptions should be secured. Think of a service provider who provides distinct descriptions for retail and wholesale clients (e.g. the wholesaler’s description would normally include a different price). This concept is similar to visibility rules in [49]. Alternatively, multiple advertisements could be generated by a service provider with access controls applied based on the type of requestor accessing the information.

Security becomes a decidedly more complex property in the context of sub-services. We propose the following questions. (1) When a client interacts with a service and authenticates it, should they also authenticate all the sub-services? Do we require security certificates that validate aggregations or compositions of sub-services? (2) How do you secure a service to stop it from being composed within another service? Securing the discovery of the service may be an alternative [16]. (3) What are the implications for a service when some sub-services require security and others don’t? (4) What happens when sub-services have differing policies with respect to client information? How do you express the security surrounding the client information to the service requestor? (5) What constitutes an infringement to a security promise? How are infringements managed (e.g. penalty payment, removal from a composition)?

4.7.2. Trust. It is easy to become very philosophical when discussing trust. As humans we use trust in a subjective manner for almost everything we do. A useful discussion of trust is offered in [33], where it is suggested that trust is a reinforcing attribute that balances perceived risk, cost and benefit. These same concerns are present in the service provision process.

Trust can be both mutual (i.e. a service provider doesn’t trust the service requestor and vice-versa) and exclusive (e.g. the service provider trusts the service requestor but the service requestor doesn’t trust the provider). A model for information flow within systems where mutual distrust is present has been offered in [34]. Service requestors largely view trust from two perspectives: whether they trust the intentions of a service provider and whether they trust the competence of a service provider.

Reputation mechanisms are an attempt to embody trust. Two such mechanisms have been offered to address the issues of misrepresentation and alteration in electronic marketplaces [52]. The implementation of reputation mechanisms may be useful but concepts from non-electronic service provision may prove useful. People tend to be satisfied that when acting within a group they will be able to increasingly trust a service provider.

The following questions arise with respect to trust in service provision. (1) How do you represent the trust of service providers or service requestors within a particular context? This question arises from a definition of reputation—“the amount of trust inspired by a particular person in a specific setting or domain of interest” [52]. (2) In a decentralised system how is knowledge relating to trust distributed, particularly changes to the perception of trust for a

party? (3) How do you trust a composition (e.g. service A is composed from sub-services X, Y and Z)? Can an external party validate a service and/or provide a level of reputation based on previous interactions? (4) What are the implications or penalties for parties that are distrustful? (5) Does access to the past performance of a service provider reduce the perceived risk of the service requestor?

4.8. *Ownership and rights*

Provision of goods usually results in a change of ownership from the service provider to the service requestor. Services don't involve a transfer of ownership. Service providers typically own the intellectual property associated with the provision process. However, service requestors *do* have a limited set of rights that are associated with a service. These rights provide a degree of control over the request and consumption of the service.

The rights available to service requestors with most services include the following. *The right to comprehend*: service requestors should be able to question the provider with the intention of better understanding a service. *The right to retract*: once an advertised service offer has been refined into a service contract, via negotiation between the service provider and the service requestor, the service requestor can choose not to request an instance of that service. The service requestor maintains the right to request the service from another service provider. *The right of premature termination*: requestors may have the ability to prematurely terminate a service. The service provider may continue provision of the service (e.g. a movie continues to play if you get up and walk out) and may choose to apply some form of penalty for partial consumption. The latter is common in the mobile phone industry where penalties apply for early termination of mobile phone plan contracts. *The right of suspension*: interrupting the delivery and therefore the consumption of a service can act as a useful method for extending the service provision process. An example of a suspension is asking the milkman to not deliver while you are on holidays. Correspondingly, *the right of resumption*: continues the delivery and consumption of a previously suspended service.

5. **The use of non-functional properties**

Non-functional properties can be used during the numerous operations of services. The service life cycle is controlled, by the service provider, from conception, to decommissioning where a service is no longer to be offered. It typically involves the definition (or creation), advertisement, invocation, and decommissioning. All these aspects form part of the general evolution of a service.

Once a service has been defined, one or more descriptions can be generated. These descriptions, sometimes referred to as advertisements or offers, are normally published with a catalogue. Matchmaking is conducted by catalogues using the search criteria provided by a requestor and the descriptions from service providers. Currently service descriptions are primarily static and are insufficient to allow detailed refinement to occur at the service catalogue. Services, and consequently their descriptions, may require modification as a result of interactions with service requestors, other service providers or their surrounding environment.

5.1. *Discovery*

As mentioned in Section 3, we consider discovery to be the process of finding candidate service providers. This does not include the refinement of the requestor's understanding of the service. Service catalogues (e.g. YellowPages) currently maintain lists of service providers categorised according to proprietary classification schemes. Non-functional properties are largely restricted to the temporal (e.g. 24×7) and spatial (e.g. address) availability, a request channel (e.g. a telephone number) and possibly a geographic region. Temporal and spatial availability for all request and delivery channels, quality of service, rights of the requestor over the service, settlement models, charging styles, security and trust are not currently provided by catalogues to requestors. Inclusion of these non-functional properties within a published description allows for more detailed refinement to occur through the service catalogue.

For example, a requestor located in Canada, wishes to discover a service that provides stock quotes from the Hong Kong Stock Exchange. The requestor wants to ensure that the following non-functional properties are addressed by the provider: (1) that the software is developed according to the ISO 9001:2000 quality standard, (2) that the request and delivery channel is the web, (3) that the settlement model is subscription-based, (4) that the charging style is by a unit of measure (i.e. time) and granularity based (i.e. monthly), (5) that payment can be made in US dollars, (6) that the information is no more than 20 minutes old, (7) that username and password security is required to access the service, (8) that they trust the service provider based on the fact that greater than 100,000 people currently utilise the service, and (9) that they have the right to terminate the service after six months with only 2 weeks notice. For non-vague service requestors, this type of service discovery is not currently possible. It is hoped that this level of description and matchmaking will reduce the need to contact providers only to discover the requestor's requirements do not match the supplied service. Publishers may wish to target specific catalogues with a more detailed description, whilst providing a high-level overview at other catalogues.

5.2. *Substitution*

Substitution uses accurate service descriptions to allow rational optimisation of sub-services within a composition. Taking two services *A* and *B* and combining them sequentially may be easy to conceptualise. Service *A* may be an electronic news report and service *B* an electronic weather report. If we try to outsource them then difficulties arise. *A* may only be offered in the USA and *B* in Chile. Pretty useless if you live in Australia; and pretty useless too if *A* is available on weekdays and *B* only on weekends. If, as virtual service builders, we want to configure such a composite service, then the non-functional properties of contributing services must be examined carefully. This discussion raises the notion of *substitutability* in the context of composition. In software engineering, there are established rules about the substitution of one function by another. These rules are captured in the approach known, not coincidentally, as programming by contract. There, we may substitute one function *F* by another *G* if *G* has weaker preconditions and stronger postconditions.

Suppose we have, at some time in the past, composed a configuration that contains A , and we encounter another potential service A' . It seems safe to assume that, if A is only available on weekdays but A' is available seven days a week then, all other things being equal, we can substitute the newer one. Thus we may anticipate a number of substitution guidelines. A' may be substituted for A provided:

- A' is cheaper than A
- A' is more spatially available than A
- A' is more temporally available than A

These rules may be compared with weakening the preconditions; for example, a service that is more geographically available has, essentially, weaker conditions attached to its use. Other properties may be associated with the concept of postcondition. For example, a service with *stronger* consumer rights may always be substituted for one with weaker obligations.

5.3. *Composition*

Composition is a way of defining a new service. Static or dynamic composition requires an accurate and detailed understanding of the services involved. Composition produces tightly-coupled integration between sub-services to ensure that value is added over the sum of the individual services. As a composer of services, discovery and substitution are integral. Discovery provides an opportunity to determine service providers that can be included in a composition, whilst substitution is useful for existing compositions where a sub-service needs to be replaced. Lets look at an example. An entity determines that they would like to compose a new service that provides hotel and car rental bookings. An appropriate hotel reservation service, and a vehicle reservation service must be found. The new service is to exhibit the following non-functional properties (1) it is to provide a single settlement model, (2) it is restricted in spatial availability accommodation and car hire in France, (3) it is restricted to service requestors from Australia, (4) the accommodation is rated as greater than 3 stars, and (5) the vehicles need to be restricted to carrying greater than 4 people. The composing entity needs to discover services that meet the specified criteria. To undertake this in a dynamic manner, sufficient functional and non-functional information must be included with its published description.

5.4. *Management*

Rich repositories of service metadata provide an opportunity for monitoring and controlling the operations that occur on that metadata (e.g. discovery, substitution, composition, execution). Existing service management architectures that support composition include Aurora and DySCo [32, 40]. We suggest that any service management architecture that aims to monitor or control service life cycle operations will need to recognise these operations by means of a rich service description language. However, such a system will need to do more. These systems may be relied on to establish that the behaviour of a service, as delivered,

is consistent with the service as specified in a contract is a highly important issue. *Conformance* may have legal consequences. How can it be demonstrated, by examination of a trace or otherwise, that a service was or is being properly delivered?

Additionally, service management repositories offer opportunities for the development of comparative tools that evaluate services “side by side” and that are capable of tracking the evolution of a particular service or type of service. As services evolve, consequently their descriptions should also reflect that metamorphosis. Evolution of a service can be the result of (a) interactions with either requestors or service composers, (b) changes to the environment that surrounds a service, (c) the need to alter the functionality, or (d) impetus from the changing constraints or non-functional properties over the service. Mechanisms that implement non-functional properties (e.g. security, trust and channels) will evolve with standards from the relevant domains. Service evolution is likely to be constrained by the existing commitments that service providers have to delivering a service. The need to administer evolving service descriptions questions the need to include expiry conditions (e.g. temporal constraints) within the description. This provides a mechanism for updating cached descriptions. A similar mechanism is provided in HTML metadata.

6. Conclusions

Whilst acknowledging the importance of service functionality, this paper has attempted to highlight the issues associated with non-functional service properties. Within this paper we have taken the view that service description is only complete once the non-functional aspects are also expressed. We believe that the ability to richly and accurately describe services has applicability in the areas of electronic service discovery, substitution, composition and management. An increased level of service property information will also facilitate more thorough decision-making by a service requestor. This paper has been motivated by the everyday services that surround us, and the ways in which we engage with them. We believe that the historical interactions, both social and economic, of commercial services offer strategic insight for the success of electronic service initiatives.

References

1. J.F. Allen, “Maintaining knowledge about temporal intervals,” *Communications of the ACM*, vol. 26, no. 11, pp. 832–843, 1983.
2. J.F. Allen, “Planning as temporal reasoning,” in J. Allen, R. Fikes, and E. Sandewall (Eds.), *Proceedings of the 2nd International Conference on Principles of Knowledge Representation and Reasoning (KR’91)*—Reprinted, Cambridge, Massachusetts, USA, 1991, Morgan Kaufmann Publishers: San Mateo, CA.
3. J.F. Allen, “Time and time again: The many ways to represent time,” *International Journal of Intelligent Systems*, vol. 6, no. 4, pp. 341–356, 1991.
4. A. Ankolekar, M. Burstein, J.R. Hobbs, O. Lassila, D.L. Martin, S.A. McIlraith, S. Narayanan, M. Paolucci, T. Payne, K. Sycara, and H. Zeng, “DAML-S: Semantic markup for web services,” in *Proceedings of the First International Semantic Web Working Symposium (SWWS): Infrastructure and Applications for the Semantic Web*, Stanford University, CA, USA, 2001, pp. 411–430.
5. C. Aurrecochea, A.T. Campbell, and L. Hauw, “A survey of QoS architectures,” *ACM/Springer-Verlag Multimedia Systems Journal (Special Issue on QoS Architecture)*, vol. 6, no. 3, pp. 138–151, 1998.

6. D. Beringer, H. Kuno, and M. Lemon, "Using WSCL in a UDDI registry 1.02: UDDI working draft best practices document," May 2001. Available from http://www.uddi.org/pubs/wsc1BPforUDDI_5_16_011.pdf, accessed on 15 Nov. 2001.
7. C. Bettini, X.S. Wang, and S. Jajodia, "A general framework and reasoning models for time granularity," in L. Chittaro, S.D. Goodwin, H.J. Hamilton, and A. Montanari (Eds.), *Proceedings of the 3rd International Workshop on Temporal Representation and Reasoning (TIME'96)*, Key West, Florida, USA, 1996, IEEE Computer Society Press, pp. 104–111.
8. W. Caelli, D. Longley, and M. Shain, *Information Security Handbook*, Macmillan Publishers: New York, USA, 1991.
9. A. Campbell, G. Coulson, and D. Hutchison, "A quality of service architecture," *ACM SIGCOMM Computer Communication Review*, vol. 24, no. 2, pp. 6–27, 1994.
10. F. Casati, S. Inicki, L. Jin, V. Krishnamoorthy, and M.-C. Shan, "Adaptive and dynamic service composition in eFlow," Technical Report HPL-2000-39, Hewlett Packard Software Technology Laboratory, Palo Alto, CA, USA, March 2000. Available from <http://www.hpl.hp.com/techreports/2000/HPL-2000-39.html>, accessed on 27 Dec. 2000.
11. H. Chen, D. Chakraborty, L. Xu, A. Joshi, and T. Finin, "Service discovery in the future electronic market," in *Workshop on Knowledge-Based Electronic Markets (KBEM'00)—Proceedings of the 17th National Conference on Artificial Intelligence (AAAI-00)*, Austin, Texas, USA, 2000, AAAI Press/MIT Press, pp. 1–6.
12. E. Christensen, G. Meredith, F. Curbera, and S. Weerawarana, "Web services description language (WSDL) 1.1. W3C Note, Ariba, Microsoft and IBM corporation," March 2001. Available from <http://www.w3.org/TR/2001/NOTE-wsdl-20010315>, accessed on 26 June 2001.
13. L. Chung, "Non-functional requirements for information system design," in R. Andersen, J.A. Bubenko Jr, and A. Sølvberg (Eds.), *Proceedings of the 3rd International Conference on Advanced Information Systems Engineering—CAISE'91*, Lecture Notes in Computer Science, Trondheim, Norway, 1991, Springer-Verlag: Berlin, pp. 5–30.
14. J. Clifford, C. Dyreson, T. Isakowitz, C.S. Jensen, and R. Snodgrass, "On the semantics of "Now" in databases," *ACM Transactions on Database Systems*, vol. 22, no. 2, pp. 171–214, 1997.
15. A.G. Cohn and S.M. Hazarika, "Qualitative spatial representation and reasoning: An overview," *Fundamenta Informaticae*, vol. 46, no. 1/2, pp. 2–32, 2001.
16. S.E. Czerwinski, B.Y. Zhao, T.D. Hodes, A.D. Joseph, and R.H. Katz, "An architecture for a secure service discovery service," in T. Imielinski and M. Steenstrup (Eds.), *Proceedings of the 5th Annual ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom '99)*, Seattle, Washington, USA, 1999, ACM: New York, pp. 24–35.
17. K. Decker, M. Williamson, and K. Sycara, "Matchmaking and brokering," in *Proceedings of the 2nd International Conference on Multi-Agent Systems (ICMAS-96)*, Kyoto, Japan, 1996, MIT Press: Cambridge, MA, pp. 1–19.
18. Discovery Universal Description and Integration (UDDI). UDDI Version 2.0 API Specification, June 2001. Available from <http://www.uddi.org/pubs/ProgrammersAPI-V2.00-Open-20010608.pdf>, accessed on 22 June 2001.
19. A. Dogac and I. Cingil, "A survey and comparison of business-to-business E-commerce frameworks," *ACM SIGecom Exchanges*, vol. 2, no. 2, pp. 16–27, 2001.
20. M. Dumas, J. O'Sullivan, M. Heravizadeh, D. Edmond, and A. ter Hofstede, "Towards a semantic framework for service description," in *Proceedings of the 9th International Federation for Information Processing (IFIP) Conference on Database Semantics—Semantic Issues in e-Commerce Systems*, Hong Kong, China, 2001, Kluwer Academic Publishers: Dordrecht.
21. A. Durante, D. Bell, L. Goldstein, J. Gustafson, and H. Kuno, "A model for the E-service marketplace," Technical Report HPL-2001-17, Hewlett Packard Laboratories, Palo Alto, CA, USA, Feb. 2000. Available from <http://www.hpl.hp.com/techreports/2000/HPL-2000-17.pdf>, access on 31 Aug. 2001.
22. A.S. Fotheringham and M. Wegener (Eds.), *Spatial Models and GIS: New Potential and New Models*, vol. 7 of GISData, Taylor and Francis: London, England, 2000.

23. F. Griffel, M. Boger, H. Weinreich, W. Lamersdorf, and M. Merz, "Electronic contracting with COSMOS—How to establish, negotiate and execute electronic contracts on the internet," in Z.M.C. Kobryn and C. Atkinson (Eds.), *Proceedings of the IEEE/OMG 2nd International Enterprise Distributed Object Computing Workshop (EDOC'98)*, La Jolla, CA, USA, 1998, IEEE Publishing.
24. International Organization for Standardization. ISO 8601:2000, 2000. Available from <http://www.iso.ch>, accessed on 07 June 2001.
25. N.R. Jennings, T.J. Norman, P. Faratin, P. O'Brien, and B. Odgers, "Autonomous agents for business process management," *Journal of Applied Artificial Intelligence*, vol. 14, no. 2, pp. 145–189, 2000.
26. C.S. Jensen, J. Clifford, S.K. Gadia, A. Segev, and R.T. Snodgrass, "A glossary of temporal database concepts," *SIGMOD Record*, vol. 21, no. 3, pp. 35–43, 1992.
27. D. Kuebler and W. Eibach, "Metering and accounting for Web services," July 2001. Available from <http://www.ibm.com/developerworks/library/ws-maws/?n=ws-7191>, accessed on 31 July 2001.
28. A. Lazcano, G. Alonso, H. Schuldt, and C. Schuler, "The WISE approach to electronic commerce," *International Journal of Computer Systems Science & Engineering*, Special issue on Flexible Workflow Technology Driving the Networked Economy, vol. 15, no. 5, pp. 343–355, 2000.
29. R.C. Lewis and B.H. Booms, "The marketing aspects of service quality," in L. Berry, G. Shostack, and G. Upah (Eds.), *Emerging Perspectives in Services Marketing*, American Marketing Association: Chicago, IL, USA, 1983, pp. 99–104.
30. F. Leymann, "Web services flow language (WSFL 1.0)," May 2001. Available from <http://www.ibm.com/software/solutions/webservices/pdf/WSFL.pdf>, accessed on 13 Aug. 2001.
31. J.K. MacKie-Mason and K. White, "Evaluating and selecting digital payment mechanisms," in G.L. Rosston and D. Waterman (Eds.), *Interconnection and the Internet. Selected Papers from the 1996 Telecommunications Policy Research Conference*, Lawrence Erlbaum: Mahwah, NJ, USA, 1997, pp. 113–134.
32. M. Marazakis, D. Papadakis, and C. Nikolaou, "The Aurora architecture for developing network-centric applications by dynamic composition of services," Technical Report TR97-0213, Institute of Computer Science, Foundation for Research and Technology, Heraklion, Greece, Dec. 1997. Available from <http://citeseer.nj.nec.com/marazakis97aurora.html>, accessed on 25 May 2001.
33. S.P. Marsh, "Formalising trust as a computational concept," PhD Thesis, University of Stirling, Stirling, Scotland, 1994.
34. A.C. Myers and B. Liskov, "A decentralized model for information flow control," in *Proceedings of the 16th ACM Symposium on Operating Systems Principles*, Saint Malo, France, 1997, ACM Press: New York, pp. 129–142.
35. W. Nagy, F. Curbera, and S. Weerawarana, "The advertisement and discovery of services (ADS) protocol for Web services," Oct. 2000. Available from <http://www.ibm.com/developerworks/library/ws-ads.html?dwzone=ws>, accessed on 13 Aug. 2001.
36. A. Parasuraman, V.A. Zeithaml, and L.L. Berry, "SERVQUAL: A multiple-item scale for measuring consumer perceptions of service quality," *Journal of Retailing*, vol. 64, no. 1, pp. 12–40, 1988.
37. J.L.A. Peiro, N. Asokan, M. Steiner, and M. Waidner, "Designing a generic payment service," *IBM Systems Journal*, vol. 37, no. 1, pp. 72–88, 1998.
38. C. Peltz, "Interacting with services on the web: A review of emerging technologies and standards for e-services," Technical report, Hewlett Packard, March 2001. Available from <http://devresource.hp.com/devresource/Docs/TechPapers/Eservices/e-services1.pdf>, accessed on 08 May 2001.
39. D. Pfoser and C.S. Jensen, "Capturing the uncertainty of moving-object representations," Technical Report CC-99-2, ChoroChronos, Athens, Greece, April 1999. Available from <http://www.dbnet.ntua.gr/~nchoros/TRs/99/2/report.ps.gz>, accessed on 05 June 2001.
40. G. Piccinelli and L. Mokrushin, "Dynamic service aggregation in electronic marketplaces," Technical Report Number 31, Hewlett Packard Laboratories, Bristol, England, Feb. 2001. Available from <http://www.hp1.hp.com/techreports/2001/HPL-2001-31.html>, accessed on 04 April 2001.
41. R.T. Rust and R.L. Oliver (Eds.), *Service Quality: New Directions in Theory and Practice*, Sage Publications: Thousand Oaks, CA, USA, 1994.

42. A. Sahai, V. Machiraju, and K. Wurster, "Managing next generation E-services," Technical Report HPL-2000-120, Hewlett-Packard Laboratories, Palo Alto, CA, USA, Sept. 2000. Available from <http://www.hp1.hp.com/techreports/2000/HPL-2000-120.html>, accessed on 27 Dec. 2000.
43. S. Saltinis, C.S. Jensen, S.T. Leutenegger, and M.A. Lopez, "Indexing the positions of continuously moving objects," Technical Report TR-44, TimeCenter, Denmark, Nov. 1999. Available from <http://www.cs.auc.dk/research/DP/tdb/TimeCenter/TimeCenterPublications/TR-44.ps.gz>, accessed on 05 June 2001.
44. T.W. Sandholm and V.R. Lesser, "Advantages of a leveled commitment contracting protocol," in Proceedings of the 13th National Conference on Artificial Intelligence (AAAI-96), vol. 1, Portland, OR, USA, 1996, AAAI Press, pp. 126-133.
45. H. Schuster, D. Georgakopoulos, A. Cichocki, and D. Baker, "Modeling and composing service-based and reference process-based multi-enterprise processes," in B. Wangler and L. Bergman (Eds.), Proceedings of the 12th International Conference on Advanced Information Systems Engineering (CAiSE'00), Lecture Notes in Computer Science, Stockholm, Sweden, 2000, Springer-Verlag: Berlin, pp. 247-263.
46. Society for Worldwide Interbank Financial Telecommunications (SWIFT), "Society for Worldwide Interbank Financial Telecommunications Web Site, 2001," Available from <http://www.swift.com/>, accessed on 08 May 2001.
47. S. Thatte, "XLANG: Web services for business process design," Web site, Microsoft Corporation, Seattle, Washington, USA, Jan. 2001. Available from http://www.gotdotnet.com/team/xml_wsspecs/xlang-c/default.htm, accessed on 22 Jan. 2002.
48. Y. Theodoridis, T.K. Sellis, A. Papadopoulos, and Y. Manolopoulos, "Specifications for efficient indexing in spatiotemporal databases," in M. Rafanelli and M. Jarke (Eds.), Proceedings of the 10th International Conference on Scientific and Statistic Database Management (SSDBM'98), Capri, Italy, 1998, IEEE Computer Society, pp. 123-132.
49. D. Trastour, C. Bartolini, and J. Gonzalez-Castillo, "A semantic web approach to service description for matchmaking of services," Technical Report Number 183, Hewlett Packard Labs, Bristol, England, July 2001. Available from <http://www.hp1.hp.com/techreports/2001/HPL-2001-183.html>, accessed on 22 Aug. 2001.
50. W3C Web Services Team, "Position papers for the world wide web consortium (W3C) workshop on web services," Technical Number HPL-2001-73, Hewlett Packard Laboratories, Palo Alto, CA, USA, April 2001. Available from <http://www.hp1.hp.com/techreports/2001/HPL-2001-73.html>, accessed on 21 May 2001.
51. J. Yang, W.-J. van den Heuvel, and M.P. Papazoglou, "Service deployment for virtual enterprises," in Proceedings of the Workshop on Information Technology for Virtual Enterprises (ITVE 2001), Gold Coast, Queensland, Australia, 2001, IEEE Computer Society, pp. 107-115.
52. G. Zacharia, A. Moukas, and P. Maes, "Collaborative reputation mechanisms in electronic marketplaces," in Proceedings of the 32nd Hawaii International Conference on System Sciences (HICSS-99)—Electronic Commerce Minitrack, Maui, Hawaii, USA, 1999, IEEE Computer Society.
53. V.A. Zeithaml and M.J. Bitner, Services Marketing, McGraw Hill: New York, USA, 1996.