Chapter 1

TOWARDS A SEMANTIC FRAMEWORK FOR SERVICE DESCRIPTION*

Marlon Dumas
Cooperative Information Systems Research Centre
Queensland University of Technology
GPO Box 2434, Brisbane QLD 4001, Australia
m.dumas@qut.edu.au

Justin O'Sullivan, Mitra Heravizadeh
GBST Holdings Pty Ltd
PO Box 1511, Milton QLD 4064, Australia
{justin.osullivan, mitra.heravizadeh}@gbst.com

David Edmond, Arthur ter Hofstede
Cooperative Information Systems Research Centre
Queensland University of Technology
GPO Box 2434, Brisbane QLD 4001, Australia
{d.edmond, a.terhofstede}@qut.edu.au

Abstract The rapid development of the Internet and of distributed computing has led to a proliferation of online service providers such as digital libraries, web information sources, electronically requestable traditional services, and even software-to-software services such as those provided by persistence and event managers. This has created a need for catalogs of services, based on description languages covering both traditional and electronic services. This paper presents a classification and a domain-independent characterisation of services, which provide a foundation for their description to potential consumers. For each of the service characteristics that we consider, we identify the range of its possible values in different settings, and when applicable, we point to alternative approaches for representing these values. The idea is that by merging

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

these individual approaches, and by mapping them into a unified notation, it is possible to design service description languages suitable for advertisement and matchmaking within specific application settings.

1. Introduction

The concept of *service* is becoming increasingly central to many areas of information technology, including digital libraries, multimedia systems, distributed computing, data management and more recently, electronic commerce. As a result, many different and often incompatible approaches to describing, managing and providing services have been developed, and there is still a lack of consensus and sometimes clear understanding about what constitutes a service.

Some recent approaches to business-to-business e-commerce [Casati et al., 2000; Schuster et al., 2000; Jennings et al., 2000], essentially view a service as a simple or a complex *task* or *activity*, executed within an organisation on behalf of a customer or another organisation. In other words, services are seen as abstractions of business processes. This abstraction is generally performed for the purpose of composition: services provided by different enterprises are composed into inter-organisational workflows, thereby leading to virtual enterprises.

Other works in the areas of middleware and database systems [Bernstein, 1996; Collet, 2000], consider a service as a set of software functionalities which facilitate the implementation of some kinds of applications. Specifically, services are seen as *software components* dedicated to a particular *aspect* of application development (e.g. transaction services, replication services, authentication services, event services). A similar definition has also been adopted in networking and telecommunications, although in these areas, a service generally involves a physical infrastructure in addition to the software component itself.

Finally, going back to the more traditional definitions found in the areas of management and marketing, a service is a product involving a performance "which results in added value in forms (such as convenience, amusement, timeliness, comfort and health) that are essentially *intangible* concerns to the first purchaser" [Zeithalm and Bitner, 1996]. Under this viewpoint, services share many characteristics with tangible products (i.e. goods): they can be bought (consumed), sold (provisioned), advertised, packaged, priced, etc. However, they fundamentally differ from goods in that they do not result in any ownership, although the right to a service can be owned. Moreover, the consumption of a service involves some kind of interaction between the consumer and the provider. As a consequence, services are generally consumed at the time they are produced [Kasper et al., 1999]. Under this definition, digital libraries, search engines, directories, and other web-based information sources, can be seen as automated service providers.

Certainly, the above conceptions of the term "service" share many commonalities. The aim of this paper is to identify some of them, in order to derive a unified view of "traditional" and "electronic" services. To this end, we study several classifications and characterisations of services, and we exploit them to lay a semantic foundation for catalogs of service offers that go beyond the current "Yellow Pages" approach to service advertisement and location.

Service description is critical to e-business application development, and has motivated many standardisation initiatives such as UDDI (Universal Description, Discovery and Integration) [Ariba Inc et al., 2000]. Backed by major companies, UDDI aims at becoming a worldwide registry for business-to-business services. As a foundation for this registry, an XML schema for describing business identities, locations and capabilities has been defined. This schema however, does not go deep into semantic aspects such as the spatial and temporal availiability of a service offer, its pricing, payment and delivery modalities, the degree of security and confidentiality of service request and consumption, etc. Instead, all these aspects are delegated to third parties. Specifically, a structure describing a service offer in an UDDI registry, contains references to one or external documentations which establish how the terms of the description should be interpreted.

Our proposal can be seen as a foundation for integrating advanced semantic aspects into service description languages such as that of UDDI. The intention is not to come up with a universal language for detailedly describing service offers. Given the broad spectrum of areas where the term "service" appears, it is unrealistic to believe that there will ever be such a universal language. Instead, we identify requirements and elements that any service description language should integrate, and we refer to relevant standards, industrial practices, and research proposals, where adequate.

The rest of the paper is structured as follows. In section 2, we discuss several classification schemes for services. Next, in section 3, we identify a set of characteristics of services which are common to a wide range of application domains and categories of services. For each of these characteristics, we describe its range of possible values, and when applicable, we outline approaches for describing these values. Finally, we provide a discussion of some related work in section 4, before drawing our conclusions in section 5.

2. Service classifications

In defining a semantic framework for service description, the interest of a classification of services is twofold: on the one hand it delineates what we mean by a service, and on the other, it structures the space of services into classes that can be more easily characterised.

Several classifications of services have been proposed in the area of services marketing and management [Kasper et al., 1999]. Lovelock's classification [Lovelock, 1983] is one of the most relevant from the viewpoint of service description. This classification is based on a set of questions that we enumerate below. In this enumeration, we have slightly modified the original formulations given in [Lovelock, 1983], so as to take into account services involving software.

- Are the actions involved in the service *tangible* or *intangible*? Internet information services such as search engines involve intangible actions, while other such as e-commerce retailing are tangible.
- Who or what is the direct recipient of the service? Is it a person, a *physical* object or a software? Reciprocally, one can ask the question about who or what is delivering the service. In this way, we obtain the following classes of services: human-to-human (hairdressing), human-to-object (equipment repair), object-to-human (vending machines), object-to-object (automatic car washer), software-to-software (event services), software-to-human (search engines) and human-to-software (software maintenance). The above classification can be further refined by noticing that in some situations, there is not a single human involved in the service delivery or consumption but rather an organisation or business (thereby leading to the term business-to-business service).
- What is the relationship between the service provider and its users? Is it a formal relationship (i.e. it requires a subscription) or not? Is the delivery of the service continuous (e.g. many services provided by operating systems) or discrete (e.g. a database query service)?
- What is the nature of demand and supply for the service? Does the demand regularly exceed the capacity (e.g. many popular search engines)? Do users have to make a reservation (e.g. some emerging bandwidth services) or are they served on a FIFO basis (e.g. memory allocation services)?
- How is the service delivered? Electronically or physically? Through a broadcast mechanism or in a point-to-point way?

This classification does not explicitly take into account at least two important issues:

- **Service automation.** In general, when the actions involved by the service are intangible, they can be partially or fully automated. This is the case for travel agencies and insurance brokers, which currently tend to fully automate their services. An extreme case of automation can be found in the area of information extraction over the Internet (e.g. search engines).
- **Service composition**. In the last decade, this issue has become crucial as business processes are being modeled through workflows, that can be connected through emerging enterprise-wide and inter-organisational workflow management systems [Casati et al., 2000; Schuster et al., 2000; Jennings et al., 2000]. As a result, services that are primarily intended to be composed with others (i.e. *intermediary* services), need to be distinguished from those which are directly consumable (i.e. *final* services). On the extreme of the automation dimension, transaction, event, and replication services, are examples of intermediary services, while travel agencies over the Internet are examples of final services.

Considering these two dimensions together, leads to a unified view of "traditional" and "electronic" services, as summarized in table 1.1.

	Fully Automated	Partially Automated	Manual
Intermediary	Transaction services Persistence services	B2B workflow-driven services	Equipment repair
Final	Web-based info sources Digital Libraries	Telephone banking E-Commerce retailing	Hairdressing Medical services

Table 1.1. Classification of services according to their degree of automation and their relationship to the consumer. The vertical axis represents the relationship of the service to its final consumers, while the horizontal axis represents its degree of automation. The contents of the cells are examples of services.

Services can also be classified according to industry branches. For instance, telephone companies' Yellow Pages rely on such classifications[1]. The Standard Industry Classification (SIC), provides an internationally recognized hierarchical classification of industries (among which services) into sectors [Investors Alliance, 1996]. Each sector in this classification is associated with a textual description and a code (e.g. 55 for hotels and restaurants, 62 for air transport). The United Nations provides another classification scheme for goods and services industries, namely UNSPSC [United Nations and Dun & Bradstreet Co, 1999]. Although

[1]Some Internet directories (e.g. Yahoo!) provide their own classifications of services into industry sectors.

the use of the above two standards is limited, they provide an invaluable foundation for service matchmaking. We note that similar widely accepted classifications are missing in the area of software services.

3. Service characteristics

To retrieve a service offer from a catalog, we consider that a user enumerates a set of characteristics, and specifies the values that (s)he is willing to accept for each of them. Given these data, the catalog system provides a list of possibly ranked candidate service offers, and enables the user to select one or several of them on the basis of both the characteristics that (s)he originally enumerated, and perhaps some others. Characterisation is therefore crucial for querying and selecting services, and needs to be taken into account during advertisement.

In order to characterise services, we need to answer the question "what does a service involve?". By systematically asking the classical W's questions, i.e. what? where? when? who? why? and how?, one may derive the following items:

What? There is an identifiable function, be that some physical activity or a computational one. By "identifiable", we mean that it is possible to give a description of the function which is understandable by the potential consumers of the service. The standard industrial classifications mentioned in the previous section provide one way of describing this function. In the general case however, the description may need to be tailored on a case by case basis in order to cater for cultural differences.

Who? Where? When? There is an identifiable trigger by which the service commences (i.e. a request), which occurs at some time and place and possibly through some channel. Once the request is processed, the *service offer* is instantiated, leading to a *service instance*, which is essentially a promise by one party (the provider), to perform a function on behalf of another party (the consumer) at some time and place and through some channel. The execution of this promise is termed *delivery*.

Why? The consumer engages to give something in exchange for the service instance (i.e. a payment), which should conform to the pricing established by the service provider. The pricing, as well as the other terms of the service delivery, can be negotiable [Jennings et al., 2000].

How? The whole process is carried out through a protocol designed to ensure some minimal guarantees and a degree of security. The execution of a service may involve human and computational activities both from the provider and from the consumer. In addition, as discussed in the previous section, the execution of a service instance may involve the instantiation of other service offers, since a service can be used as part of others (composability).

On the basis of the above enumeration, we can identify the following characteristics of a service offer: provider, availability, channel, pricing, payment, security, quality of service, and reputation. These characteristics are transversal to the categories of services discussed in the previous section, although their range of values may differ from one category to another (e.g. whether the service is fully automated or semi-automated, or whether it is software-to-software or software-to-human). For this reason, they can be used as a common framework for querying a catalog of heterogeneous service offers, shortlisting the query answers, selecting an offer, and requesting the service. In the sequel, we examine each of these characteristics, except the "provider" one, for which the syntax and semantics are straightforward. Although we consider each characteristic independently of the others, it should be noted that in practice they are often correlated (e.g. the pricing may depend on the security or the quality of service).

3.1. Temporal and spatial availability

Before defining temporal and spatial availability, it is important to distinguish the time and the place of a service request (i.e. booking), from the time and the place of its delivery. To this end, we define the *request time* (resp. *request location*) as the moment (resp. *place*) at which a given customer requests the service. Similarly, the *delivery time* and *location* refer to the moment and place when/where an instance of the service is consumed.

With these definitions, temporal and spatial availability may be modeled as a set of restrictions over the above four parameters. These restrictions may concern each of the four parameters individually, or they may express some inter-relationship between them. In the former case, the constraints over the time parameters can be expressed as a set of instants, while the constraints over the locations can be formulated as a set of points. The latter case can be further decomposed into two: either the inter-relationship concerns times and locations separately (e.g. the request must be performed between 3 and 5 days prior to the delivery), or there is an inter-relationship between a time and a location (e.g. the service is delivered at a given location for some period of time, and at another one after this period). The first situation can be captured by introducing temporal and spatial constraints separately (e.g.

$request\_time \leq delivery\_time - 3\ days$)[2]. The second case requires one to express time and space in a single reference system, thereby making spatio-temporal objects an interesting candidate representation for the availability of a service offer, as discussed below.

At a concrete level, a set of instants can be represented as a period (e.g. a promotional service is available between 1/1/2001 and 31/3/2001), or as a sequence of disjoint and non-contiguous periods (e.g. a guided tour which is available during the period [1/1/01..31/3/01] and [16/4/01..30/6/01]). In many realistic scenarios, the set of availability instants of a service (whether regarding the request or the delivery) exhibit some kind of periodicity (e.g. the opening hours of a bank). In such situations, a representation based on "calendars", such as those proposed in [Leban et al., 1986] and [Chandra et al., 1994], can be far more adequate. These formalisms support the expression of sets of instants such as *"8am through 4pm of every working day between 1/1/2001 and 31/3/2001"*. In any case, each of the instant literals involved in the representation of a set of instants, can be expressed in several formats. The ISO standard 8601:1988 is intended to serve as a reference format for dates and times, but its use is quite limited. Extensible date and time format systems such as those proposed in the TSQL2 language [Snodgrass, 1995] should therefore be considered instead.

On the other hand, the issue of representing sets of points has been extensively addressed by the spatial database and the spatial reasoning communities [Rigaux et al., 2001]. Although many alternative representations have been studied, simple vectorial representations are the most commonly used, especially within geographical information systems. We can therefore safely adopt the point of view that the spatial availability of a service is expressed as a point, a set of points, a polygon, or a set of polygons. Alternatively, a spatial logical identifier (e.g. the name of a city or a suburb) can be used instead of the actual spatial location. In this case, either the description of this reference is based on an agreed-upon format (e.g. street names, postal codes, and country codes[3]), or a reference to a documentation must be provided, so that the user can interpret this identifier. This reference can be modeled using the concept of TModel introduced in UDDI [Ariba Inc et al., 2000]. Roughly speaking, a TModel is a reference to a resource (e.g. a web site) that provides the documentation for understanding a term within a service description. The disadvantage of using TModels is that in the general

[2]Notice that by definition, the request time is constrained to precede the delivery time.

[3]For country codes, see the ISO 3166 standard. For a detailed approach to "address description", see the xCBL documentation [Commerce One Inc., 2000].

case, the documentation is not in a format which allows a software to exploit it. This is the case if the documentation is an image containing a map of the location and its surroundings.

The issue of representing spatio-temporal objects has been extensively addressed in the area of spatio-temporal databases (see e.g. [Erwig et al., 1999]). However, the existing approaches in this area do not handle situations where temporal periodicity is involved. For this reason, we prefer a representation of spatio-temporal points based on pairs composed of a spatial region and a set of instants. For instance, the spatio-temporal availability of an opera performance can be expressed as follows[4]:

- ⟨Queensland Performing Arts Centre, TModelQPC⟩: Saturdays and Sundays between 15/2/01 and 28/2/01. (TModelQPC is a reference to a TModel.)
- ⟨Sydney Opera House, TModelSOH⟩: every day except Mondays, between 1/4/01 and 15/6/01. (TModelSOH is a reference to a TModel.)

Many services are requestable or delivered "at arms length" through some electronic channel as discussed in the next paragraph.

3.2. Request and delivery channels

With the introduction of the Internet and of new communication devices (mobile phones, pagers, etc.), there has been an increase in the number of request and delivery channels available to consumers. This has not only increased the flexibility of the service offerings, but has also pushed the providers to ensure the continual upgrade of their service.

A *channel* is the means by which a user requests a service or receives the resultant output from a service. These are referred to as the request and delivery channels respectively.

To further illustrate the concepts of request and delivery channels we consider a concrete example. A day trader utilising the services of a brokerage house may place trades using either of the following methods: a Web-based online trading system, an Interactive Voice Response (IVR), a Wireless Application Protocol (WAP) [WAP Forum, 2001] a mobile phone or personal digital assistant (PDA), and a standard telephone service (e.g. calling an advisor). These means of access are called *request channels*. On the other hand, the brokerage house may offer a notification service for price changes (e.g. the value of stock MSFT on

[4]For the sake of simplicity, we do not introduce any concrete notation for sets of instants. Instead, we refer the reader to [Leban et al., 1986] and [Chandra et al., 1994].

the NASDAQ exchange reaching the price $'x.xxx'), such that the alerts can be configured for delivery through several channels: email, Short Message Service (SMS), pager and facsimile. These means are called the *delivery channels*.

It should be noted that electronic delivery channels are primarily relevant to *information services* (both addressed to persons or to softwares). The delivery channel of services involving a physical object delivered at arms length (see section 2), is necessarily a transportation means (e.g. postal mail, cargo, etc.).

Delivery channels may be *broadcast* mechanisms whereby all relevant information is "pushed" to the requesting user. Security of the request and delivery channels may be required. We address security in a separate subsection below.

A syntax for request and delivery channels should take into account the following aspects:

- Availability: Depending on the physical or electronic nature of the service provision, the delivery location may be represented as a mail address (e.g. post office box), an Internet address (e-mail, URI, or IP address including the port number), a telephone or facsimile number (including area and country codes), or any other telecommunication resource address. The UDDI proposal defines a syntax for such descriptors. In the case where the request is performed through a software, a description of the protocol for requesting the service provider should accompany the location data. An approach for describing this protocol is provided by the Web Services Description Language (WSDL) [Christensen et al., 2000].
- Protocol: The method used for communicating with the end point (e.g. in a financial scenario this may be the Financial Information Exchange (FIX) protocol).
- Operations: Functions that can be invoked by the service instantiator in a request scenario or by the service provider in a delivery scenario.
- Security Model (Applicable to electronically delivered services): Approach used for information transmission during the service delivery (e.g. SSL).

3.3. Payment and pricing

Payment is the business process defined by the service provider for collecting the price of the service from the consumer. Payment can be conducted in single or multiple stages (i.e. installments), using various mediums (e.g. direct cash exchange, credit or debit card, cheque, direct debit, etc.) and at different stages within the service provision process. Payment can be done before delivery, at delivery, after delivery, or any combination of the above.

In some situations the obligation of payment is waived. For instance, the use of a freeware is not subject to a payment unless it is used for commercial purposes. This is different from the situations where a service is free when it is accompanied by another service (e.g. a mobile phone is provided for free if the customer agrees to a 1-year contract with the telecommunications provider). In this case, the conjunction of the "free" and the "paying" services form a *package*, which consitutes a service per se.

Pricing is normally a function of the service provider recouping wholesale cost and adding a profit margin, or a market environment displaying normal supply and demand characteristics (e.g. a stock market). Pricing for a service is largely at the discretion of the service provider and as such, we consider a service to have a nominal price. It should be noted that in some domains the existence of an organised body (i.e. a cartel) is used to define the price of services. Consumers wishing to reduce the cost of service provision can sometimes form consumer groups to achieve economies of scale. We term these *cooperatives*.

The following are elements that a notation for pricing and payment should include:

- Price: Charge for the service being provided (normally represented as a number with optional decimal points).
- Settlement Currency: Standard representation for the currency utilised during the settlement process (e.g. ISO 4217:1995 standard).
- Settlement Date/Time: Date and time represented using a defined standard (e.g. ISO 8601:1988) expressed with respect to a coordinated universal time.
- Payment Schedule: When the payment involves several transactions (e.g. a lease), a matrix is used to represent the combination of amount, settlement date/time and pre-condition values for each of them. Each row in the matrix is numbered, and these numbers are used as references in the body of the preconditions.
- Payment System: A description or a reference to a mode of payment (i.e. Cash, Credit Card, Cheque). Many electronic payment models have emerged recently, and to the best of our knowledge, there are currently no comprehensive standards for identifying them. [Mahony et al., 1997] provides a survey of this area.

- Payment Channel: Method used for conducting the payment (e.g. Closed Network, Internet, Email, Post and Phone). This is different to the service delivery channel discussed previously.
- Security Model: Approach used to ensure a degree of security over the payment channel (e.g. SSL).
- Beneficiary: Identifier of the entity to which the payment is addressed, which can be different from the ID of the service provider. Examples include a UDDI Business ID, a Market/Exchange ID (e.g. NASDAQ, eBay), or the name of a company.
- Penalty Cost Schedule: Similar to the payment schedule, this is used to define the penalties for not completing the payment.

Price and payment can be tightly related. For instance, the price of a service can depend on the time of payment and/or its division into installments. Pricing and payment are tightly linked to the business model of the service provider. Characterising pricing and payment is therefore equivalent to characterising business models which is a quite complex problem (see e.g. [Rappa, 2000] for a discussion on this issue in the context of e-commerce). The details provided above represent just a summarised view.

3.4. Security

Security of a service, or a part of a service should be configurable by both the service provider and the service consumer. Security is usually defined using four dimensions [W. Caelli, 1991]: integrity (ensuring information is not altered), confidentiality (cryptographic techniques applied to the information), non-repudiation (ensures receiving parties cannot renege on the receipt of the information) and authentication (confirm the intended recipient and identify the originator). These dimensions introduce a level of trust that can strengthen the reputation of a service.

Specific providers may impose a high level of security when delivering services (e.g. in the banking and financial area). Numerous mechanisms may be used within specific domains to ensure security. Banks for instance secure the access to their services through magnetic cards, passbooks, pin numbers, and for services such as Internet banking, customer identifiers and passwords. In addition, customers are normally required to present identification (e.g. drivers license, passport) prior to accessing the bank's services over the counter. A description of the security mechanisms attached to a service may be part of a service offer, and in some situations, it may be used as a selection criterium.

The request, payment and delivery of a service usually involve some information exchange between the consumer and the provider. The confidentiality of this information therefore becomes an issue. Several levels of confidentiality can be identified, ranging from not revealing or making accessible this information to third parties, to partially restricting the access to the data even to one or several of the entities involved in the provisioning of the service.

In the setting of web-accessible services, confidentiality can be achieved by using standard encryption mechanisms during the transmission of the data (e.g. using the Secure Socket Layer (SSL) protocol), by carefully controlling the access to the databases that store the customer's data, and by cross-checking some information whenever the customer contacts the provider (e.g. asking the date of birth). The mechanism(s) used to ensure confidentiality can be regarded as a parameter of a service.

Elements of a notation for describing service security include:

- Certificates: Both personal and corporate identification of service parties. Certificates can also be used between applications in an insecure environment.
- Signatures: Normally attached to a document or message with the signer's public key signature. This ensures the non-repudiation of the item.
- Encryption: Algorithms used to ensure confidentiality of information (e.g. RC4). This should include a representation of the strength of the technique (e.g. 128-bit).
- Data Integrity: Use of message digests and hashes to ensure messages have not been altered during transit.
- Key Management and Storage: Appropriate techniques applied to ensure that keys are managed and maintained for both service providers and service requesters.
- Auditing Level: An important aspect of the security for any service is the ability to trace the interactions that have occurred during the execution of its instances. Auditing should not only track the parties involved, but the service being requested, request/delivery channel, temporal aspects, etc. An enumeration of the aspects of the service executions that are traced may be part of a published service description.

3.5. Quality of service

We believe that Quality of Service (QoS) is a domain-specific characteristic that has two dimensions. Firstly, the service consumer's expectations of the service being requested. These expectations can be derived from previous experiences that consumers have had with the service.

Secondly, QoS can relate to the level of commitment that the service provider has to completing the service request. This dimension represents a *warranty* that is provided to the consumer. This type of QoS may be formalised using *Service Level Agreements* (SLA). These are binding contracts entered into by the service provider and the service instantiator. Failure to provide the service at the agreed levels normally introduces some form of penalty payment. SLAs can be used to ensure quality at a course-grained level (i.e. the entire service) or components of the service (i.e. pricing and payment, temporal and spatial availability).

Commitment to a service can be bound into the contracting protocol [Sandholm and Lesser, 1996]. This approach offers a means of decommiting from a transaction, assuming that an associated penalty is paid. QoS is an auditable aspect of most services.

Elements of a syntax for QoS include the following:

- Accessibility: A measure of the access/uptime of the service.
- Performance: A measure of the speed of service execution.
- Conformance: The probability that a service provider's service level agreement is fulfilled in a particular situation.
- Guarantee: A de-commitment penalty.
- Reliability: A measure of the probability of success of the transactions involved in the service provisioning. This aspect is orthogonal to the above: the success may be defined with respect to the availability, the ease of use, the performance of execution, and the conformance to a service level agreement. In the case where this probability is calculated by considering past transactions, this value can be seen as a measure of reputation.

3.6. Reputation

This characteristic of services encompasses numerous factors, including past experience of consumers with the service, brand awareness through advertising, and adherence to a quality management standard. Past experience can be measured in several ways (e.g. as a rate of "sucessful" service executions). This rating can be provided by a third party, or obtained through referral systems involving previous consumers. Amazon.com and other online book sellers, request reviews from book purchasers in an attempt to assist their users with the product selection process. Adherence to a quality standard is more difficult to measure, although certifications address in some way this problem.

A notation for reputation should include the following elements:

- Brand Identifier.
- Industry Standard(s): General well-known standards that the service conforms to (e.g. ISO 9000 series, and in particular ISO 9001:2000 certification).
- Mediated Ratings: Provided by independent third-parties (e.g. industry chambers), these ratings reflect previous executions of the service. Ratings may take into consideration factors such as customer satisfaction, and conformance to SLA. Each rating is accompanied by a reference to the mediator that provides it.
- Referrals: Ratings and evaluations provided by the consumers.

4. Related work

The concept of service has been studied in many areas, including marketing, business management, workflow management, digital libraries, networking, and distributed computing. For space reasons, we limit our discussion on related works to those directly concerned with the scope of this paper, that is, service catalogs and their corresponding service description languages.

4.1. Product and service catalogs

There are numerous approaches to represent and query product catalogs. Although some of these approaches can be applied to services, they do not take into account their specificities, such as the temporal and spatial availability, the delivery channel, the pricing and related contractual issues, etc. Some electronic Yellow Pages capture the availability and, to some extent, the channel, by relying on proprietary user interfaces and representation structures for expressing and evaluating user queries over these characteristics. The implementation of these features is facilitated by the fact that Yellow Pages are geographically restricted.

The recent UDDI initiative referenced throughout this paper, has the ambition to become a worldwide registry for business-to-business services. It relies on an XML schema for describing the identities, contact details, and services provided by businesses. This schema delegates advanced semantic issues such as categorisation, to third party models, by introducing the concept of TModel : an annotated reference to an external documentation. The classification and characterisation effort reported in this paper can be used as a common framework for expressing TModels and their associated documentation.

Information exchange between catalogs is currently restricted due to their heterogeneity. Given the inherent customisability of most services, this heterogeneity is likely to be a central issue in the context of service catalogs. [Ng et al., 2000] considers two possible approaches to address this issue: standardisation and integration. Whilst standardisation provides a common vocabulary for undertaking information exchange between service catalogs, it is presently limited by the depth of existing characterisations and classifications of services. [Investors Alliance, 1996; United Nations and Dun & Bradstreet Co, 1999] provide hierarchical classification schemes that attempt to define global standards for the identification of goods and services. Unfortunately, these classifications only capture industrial sectors. Our proposal complements these standards by synthesising classification and characterisation schemes which are transversal to industry sectors. The integration of service catalogs on the other hand, is troubled by the need to establish mappings between them, which requires the identification of a common semantic framework for service description. Our classification and characterisation effort is precisely a first step towards this framework.

4.2. Service description languages

Perhaps one of the closest works to ours is the service description framework of the Open Service Model [Merz et al., 1997]. This framework identifies properties of service offers that are relevant for their indexation within catalogs. Specifically, the following properties of service offers are identified:

- The identifier of the service provider and the reference of the service offer within the catalog of that provider.
- The URL of the interface to the service (in the case of electronically requestable services).
- Price information (including currency).
- The initial and final availability dates.
- The service semantics and commercial conditions.

These properties are encompassed by our characterisation of services.

The XML Common Business Library (xCBL) [Commerce One Inc., 2000] provides a set of schemas for business-to-business (B2B) document exchange, in the form of XML DTDs and SOX schemas [W3C, 2000]. Based on previous Electronic Data Interchange (EDI) standards, xCBL is built upon a set of *document schema components*, corresponding to situations that are considered to occur frequently in B2B interactions: direct and indirect procurement, planning, auctions, purchase orders, invoicing, and payment. The schemas for purchase orders, invoicing and payment, capture some of the characteristics of services discussed in this paper. In addition, xCBL provides pieces of schemas (called "building blocks") corresponding to fields such as postal addresses, dates, currencies, and industry branches, which could be easily reused within a service description language. Hence, even if xCBL would need to be extended in order to accommodate all the characteristics of services discussed in this paper, its current version does show that a standardisation approach to service description is indeed feasible, and could lead to increased automation in B2B interactions.

The Web Service Description Language (WSDL) [Christensen et al., 2000] allows a developer to describe how a web-based software-to-software service can be invoked, but it does not consider its capabilities nor its contracting conditions (e.g. availability, price, and payment model). In a way, WSDL's scope of applicability is similar to that of component interface definition languages [Szyperski, 1998] such as CORBA's IDL.

In fact, the boarder between software-to-software services and software components is not very clear. Perhaps the main differences rely on their users (or more aptly, their markets). Components are developed for, and used by programmers and software developers, while services can be deployed for a much wider community. In this respect, the remark on p. 340 of [Szyperski, 1998] that "components are not necessarily at a level of granularity that makes any sense to end users" is of interest. Services typically are at a level of granularity meaningful to end users. In addition, services may involve human tasks, which makes them interesting for abstracting functionalities that may be either purely computational or not, depending on the invocation context.

Service composition platforms such as eFlow [Casati et al., 2000] and CMI [Schuster et al., 2000], provide languages for expressing control and data flow among electronically requestable services involved in an inter-organisational workflow. These proposals are complementary to ours as they do not address the issue of describing atomic services.

Another family of proposals complementary to ours is that of agent capability description languages [Sycara et al., 1999]. These languages support the description of the context of usability and outcome of the services provided by an agent, and are designed to be used by matchmaking agents (i.e. agents whose role is to locate other agents).

5. Conclusion and future work

Based on an extensive analysis of existing works in the areas of services marketing, virtual enterprises, and software services, we have developed a classification and a domain-independent characterisation of services, which together provide a foundation for their description to potential consumers. In a context where the notion of service has become ubiquitous, and where the services provided by emerging and established organisations are becoming increasingly manifold, complex and dynamic, it is clear that this kind of effort is essential for at least two purposes:

- Designing languages for describing entries within catalogs of services, and queries over these catalogs.
- Establishing a formal background for reasoning about services. In particular, it should be possible from the description of two services, to determine if they can be composed and to derive some properties of their composition.

The work reported in this paper is just a first step towards these objectives. The characterisation that we have proposed should be further refined. For instance, languages for describing the interaction between the provider and the consumer during the delivery process need to be designed. Furthermore, it should be possible to describe the outcome of a service execution, i.e. the "state" to which it leads, in terms of its preconditions. This effort should build on existing works in the areas of components and agents capabilities description languages.

On the long term, we expect that this work will lead to an extensible service advertisement language. Extensibility is a key requirement for such a language, since it should accommodate domain-specific characteristics and ontologies.

Another research avenue that we plan to pursue, is that of *service specialisation* [Malone et al., 1999]. Service specialisation underlies any efforts of service classification, which in turn are essential for structuring any catalog of service offers. While the SIC and the UNSPC classification schemes (see section 2) rely on purely functional aspects (i.e. what is the functionality behind a service), one could imagine classification schemes based on any other form of specialisation. For example, the class of services "5-star accommodation" can be seen as a specialisation of the class of services "accommodation", in which one of the characteristics (i.e. the QoS) is constrained. The issue of service specialisation is also crucial for customisation and for delegation, i.e. determining whether a service offer can be replaced by another one.

References

Ariba Inc, Microsoft Co, and IBM Co (2000). Universal Description, Discovery and Integration of business for the web. `http://www.uddi.org`.

Bernstein, P. (1996). Middleware: A model for distributed systems services. *Communications of the ACM*, 39(2):86–98.

Casati, F., Ilnicki, S., Jin, L.-J., Krishnamoorthy, V., and Shan, M.-C. (2000). Adaptive and dynamic service composition in eFlow. In *Proc. of the Int. Conference on Advanced Information Systems Engineering (CAiSE)*, Stockholm, Sweden. Springer Verlag.

Chandra, R., Segev, A., and Stonebraker, M. (1994). Implementing calendars and temporal rules in the next generation databases. In *Proc. of the IEEE Int. Conference on Data Engineering (ICDE)*, Houston TX, USA. IEEE Press.

Christensen, E., Curbera, F., Meredith, G., and Weerawarana, S. (2000). Web Services Description Language (WSDL). `http://msdn.microsoft.com/xml/general/wsdl.asp`.

Collet, C. (2000). The NODS project: Networked Open Database Services. In *Proc. of the Int. Symposium on Objects and Databases*, Sophia Antipolis, France. Springer Verlag.

Commerce One Inc. (2000). XML Common Business Library (xCBL). `http://www.xcbl.org`.

Erwig, M., Guting, R., Schneider, M., and Vazirgiannis, M. (1999). Spatio-temporal data types: An approach to modeling and querying moving objects in databases. *GeoInformatica*, 3(3).

Investors Alliance (1996). Standard Industrial Classification. `http://powerinvestor.com/SIC/`.

Jennings, N., Norman, T., Faratin, P., O'Brien, P., and Odgers, B. (2000). Autonomous agents for business process management. *Journal of Applied Artificial Intelligence*, 14(2):145–189.

Kasper, H., Helsdingen, P. v., and Vries, W. d. (1999). *Services Marketing Management: an International Perspective*. John Wiley & Sons, Chichester, UK.

Leban, B., McDonald, D., and Forster, D. (1986). A representation for collections of temporal intervals. In *Proc. of the 5th National Conference on Artificial Intelligence (AAAI)*, Philadelphia PA, USA. Morgan Kaufmann.

Lovelock, C. (1983). Classifying services to gain strategic marketing insights. *Journal of Marketing*, 47:9–20.

Mahony, D., Peirce, M., and Tewari, H. (1997). *Electronic Payment Systems*. Artech House Publishers, Norwood MA, USA.

Malone, T., Crowston, K., Lee, J., Pentland, B., Dellarocas, C., Wyner, G., Quimby, J., Osborn, C., Bernstein, A., Herman, G., Klein, M., and O'Donnell, E. (1999). Tools for inventing organisations: Toward a handbook for organisational processes. *Management Science*, 45(3).

Merz, M., Witthaut, M., and McConnel, S. (1997). OSM D8: Catalogue and service architecture. http://osm-www.informatik.uni-hamburg.de/osm-www/public/docs.

Ng, W., Yan, G., and Lim, E.-P. (2000). Heterogeneous product description in electronic commerce. *ACM SIGeCom Exchanges*, 1(1):7 – 13.

Rappa, M. (2000). Business models on the web. http://ecommerce.ncsu.edu/business_models.html.

Rigaux, P., Scholl, M., and Voisard, A. (2001). *Spatial Database Management Systems: Applications to GIS*. Morgan Kaufmann.

Sandholm, T. W. and Lesser, V. R. (1996). Advantages of a leveled commitment contracting protocol. In *Proceedings of the Thirteenth National Conference on Artificial Intelligence (AAAI-96)*, pages 126 – 133, Portland, OR, USA. AAAI Press.

Schuster, H., Georgakopoulos, D., Cichocki, A., and Baker, D. (2000). Modeling and composing service-based and reference process-based multi-enterprise processes. In *Proc. of the Int. Conference on Advanced Information Systems Engineering (CAiSE)*, Stockholm, Sweden. Springer Verlag.

Snodgrass, R. T., editor (1995). *The TSQL2 temporal query language*. Kluwer Academic Publishers.

Sycara, K., Klusch, M., Widoff, S., and Lu, J. (1999). Dynamic service matchmaking among agents in open information environments. *SIGMOD Record*, 28(1):47–53.

Szyperski, C. (1998). *Component software: Beyond object-oriented programming*. Addison-Wesley, Reading MA, USA.

United Nations and Dun & Bradstreet Co (1999). United Nations Standard Products and Services Codes. http://eccma.org/unspsc.

W. Caelli, D. Longley, M. S. (1991). *Information Security Handbook*. Macmillan Publishers.

W3C (1997-2000). Schema for Object-Oriented XML. http://www13.w3.org/TR/NOTE-SOX.

WAP Forum (1999-2001). Wireless application protocol. http://www.wapforum.org/.

Zeithalm, V. and Bitner, M. (1996). *Services Marketing*. McGraw Hill, New York, NY USA.