

# 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*

David Edmond, Arthur ter Hofstede

*Cooperative Information Systems Research Centre*

*Queensland University of Technology*

*GPO Box 2434, Brisbane QLD 4001, Australia*

**Abstract** The inexpensive and global connectivity provided by the Internet has triggered a wave of interest in providing service-oriented electronic access to commercial activities. This pressure has led, in turn, to a need for accurate service description, so that we may advertise, locate, analyse and compare services. In this paper, we classify services by the context in which they are used. Next, we characterise both conventional and electronic services according to a range of domain independent attributes including price, payment method and availability. We examine possible representations for each of these service dimensions. By integrating these representations into a unified service description language, we hope to provide a means to lubricate the electronic services marketplace.

---

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

## 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. Consequently, a clear understanding and consensus about what constitutes a service has not been reached.

Some recent approaches to business-to-business e-commerce [Casati et al., 2000; Schuster et al., 2000; Jennings et al., 2000], view a service as a simple or a complex *task* or *activity*, executed within an organisation on behalf of a customer or organisation. In other words, services are seen as abstractions of business processes. This abstraction is generally performed for the purpose of composition: where 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], 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). A similar definition has also been adopted in networking and telecommunications.

Finally, more traditional management and marketing definitions view a service as 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” [Zeithaml 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.

Service description is critical to e-business application development, and has motivated standardisation initiatives such as UDDI (Universal Description, Discovery and Integration) [Ariba Inc et al., 2000]. UDDI

lacks semantic aspects such as spatial and temporal availability, the degree of security, etc. Instead, all these aspects are delegated to third parties. Our proposal can be seen as a foundation for integrating advanced semantic aspects into description languages such as that of UDDI. The intention is not to come up with a detailed universal language for describing service offers. Instead, we identify requirements and elements that any service description language should integrate.

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. For each characteristic, 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, our interest in a classification of services is twofold: it delineates what we mean by a service, and it structures 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 particularly relevant from a service description viewpoint. This classification is based on a set of questions that we enumerate below. We have slightly modified the original formulations so as to take into account services involving software.

- Who or what is the direct recipient of the service? Is it a person, a *physical* object or 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).
- 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. 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, subscription based or via a point-to-point mechanism?

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

- **Service automation.** In general, when the actions of the service are intangible, they can be partially or fully automated. This is the case for travel agencies and insurance brokers.
- **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 for composition with others (i.e. *intermediary* services), need to be distinguished from those that are directly consumable (i.e. *final* services).

Considering these two dimensions together, leads to a unified view of “traditional” and “electronic” services, as summarised in table 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.* Classification of services according to their degree of automation and their relationship to the consumer. The rows represent the relationship of the service to its final consumers, while the columns represent the degree of automation. The contents of the cells are examples of services.

Several industry standards also exist for classifying services. The Standard Industry Classification (SIC), provides an internationally recognized hierarchical classification of industries into sectors (including service industries) [Investors Alliance, 1996]. The United Nations provides another classification scheme for goods and services, namely UN-SPSC [United Nations and Dun & Bradstreet Co, 1999]. Although the use of the above 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 catalogue, 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 this data, the catalogue 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 systematically ask the classical W's questions, i.e. who? what? where? when? how? and why?:

**What?** There is an identifiable functionality, be it a physical or computational activity. By "identifiable", we mean that it is possible to describe the function which is in-turn comprehended by the potential service consumers. The standard industrial classifications mentioned previously provide one way of describing this function. In general however, the description may need to be tailored on a case by case basis.

**Who? Where? When?** There is an identifiable trigger by which the service commences (i.e. a request), which occurs at a time and place, and via a 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 (e.g. 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 the execution of a service instance may involve the instantiation of other service offers, since a service can be used as part of another service (composability).

Using these enumerations, 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 catalogue of heterogeneous service offers, shortlisting the candidates, selecting an offer, and requesting the service. We examine each of these characteristics, except the "provider", for which the syntax and semantics

are straightforward. Although we consider each characteristic independently, it should be noted that in practice they are often correlated (e.g. the pricing may depend on 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.  $request\_time \leq delivery\_time - 3\ days$ )<sup>1</sup>. The second case requires the expression of 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. an advertised service is available between 1/11/2001 and 31/3/2002), or as a sequence of disjoint and non-contiguous periods (e.g. a guided tour which is available during the period [1/11/01 to 31/3/02] and [16/4/02 to 30/6/02]). 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

---

<sup>1</sup>Notice that by definition, the request time is constrained to precede the delivery time.

such as “8am through 4pm of every working day between 1/11/2001 and 31/3/2002”. In any case, each of the instant literals involved in the representation of a set of instants, can be expressed in several formats (e.g. ISO 8601:2000). Extensible date and time format systems such as those proposed in the TSQL2 language [Snodgrass, 1995] should be considered.

On the other hand, the issue of representing sets of points has been extensively addressed by the spatial database and the spatial reasoning communities. 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<sup>2</sup>), or a reference to a document 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 generally the documentation is not in a format which allows software to exploit it (e.g. 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: ⟨Sydney Opera House, TModelSOH⟩: 9am - 5pm daily except Mondays, between 1/4/02 and 15/6/02. (TModelSOH is a reference to a TModel.)<sup>3</sup>

Many services are requestable or delivered “at arms length” through some electronic channel as discussed in the next section.

---

<sup>2</sup>For country codes, see the ISO 3166 standard. For a detailed approach to “address description”, see the xCBL documentation [Commerce One Inc., 2000].

<sup>3</sup>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].

### 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] enabled mobile phone, or via the telephone (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 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) or pager. These channels 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 software). 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(s). 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 aspects such as location (physical or electronic), protocol, specific operations, temporal availability and the security model.

### 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 (prior to 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 freeware is not subject to payment unless 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 constitutes a service per se.

Pricing is normally a function of the service provider recouping whole-sale 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. 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 (e.g. cooperatives) to achieve economies of scale.

Price and payment are closely 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 following are elements of a notation for pricing and payment: price, currency, payment schedule, payment system, payment channel, the security model, the beneficiary and the penalty cost schedule.

### 3.4. Security

Security of a service, or a part there-of should be configurable by either the service provider or the service consumer. Security is usually defined along four dimensions [Caelli et al., 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). Banks for instance secure the access to their services through magnetic cards, pin numbers, customer identifiers and/or passwords. In some situations, a description

of the security mechanisms of a service may be used as an important selection criterion. Several levels of confidentiality can be identified for the bi-directional exchange of information between the consumer and the provider. These range from not revealing or making accessible this information to third parties, to partially restricting the access to this information to entities involved in the provisioning of the service.

With web services, confidentiality is commonly achieved by using standard encryption mechanisms during the transmission of the data (e.g. Secure Socket Layer (SSL) protocol). 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, digital signatures, encryption algorithms, data integrity mechanisms, key management and storage, and auditing levels.

### 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 consumer experiences. 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. 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). Failure to provide the service at the agreed levels normally introduces some form of penalty payment. Commitment to a service can be bound into the contracting protocol [Sandholm and Lesser, 1996]. This approach offers a means of de-committing from a transaction, assuming that an associated penalty is paid.

Elements of a notation for QoS include accessibility (a measure of the uptime of the service), performance (a measure of the speed of service execution), confirmance (the probability that the service provider's service level agreement is fulfilled), guarantee (de-commitment penalty) and reliability (a measure of the success of the transactions involved in the service provisioning).

### 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, for example as a rate of “successful” service executions. Such a rating can also be provided by a third party, or obtained through referral systems involving previous consumers. For example, 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 (e.g. ISO 9000 series, and in particular ISO 9001:2000 certification) is more difficult to measure, although certifications address in some way this problem.

## **4. Related work**

The concept of service has been studied in many areas, including marketing, business management, workflows, digital libraries, networking, and distributed computing. We limit our discussion on related work to those directly concerned with the scope of this paper, that is, service catalogues and their corresponding service description languages.

### **4.1. Product and service catalogues**

There are numerous approaches to represent and query product catalogues. 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, and the pricing, etc. Some catalogues (e.g. Yellow Pages) rely upon proprietary representation structures for expressing service characteristics.

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 catalogues is currently restricted due to their heterogeneity. [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 catalogues, it is presently limited by the depth of existing characterisations and classifications of services. [Investors Alliance, 1996; United Nations and Dun & Bradstreet Co, 1999] pro-

vide 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 catalogues on the other hand, is troubled by the need to establish mappings between them, which requires the identification of a common semantic framework. Our classification and characterisation effort is precisely a first step towards this framework.

## 4.2. Service description languages

Perhaps the closest work 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 catalogues. Specifically, the following properties of service offers are identified: service provider, the service offer identifier, the URL to the interface of the service, price information, initial and final availability dates and the service semantics (expressed in plain text). 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. In addition, xCBL provides pieces of schemas (called “building blocks”) corresponding to fields such as postal addresses, dates, which could be easily reused within a service description language. Although needing extension to support all service characteristics, xCBL does show that a standardisation approach to service description is indeed feasible.

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). 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 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.

DAML-S [DAML Services Coalition, 2001] is an ongoing effort to define a markup language for describing services, as well as user preferences and constraints over the use of services. The language is intended to be used by software agents for service discovery and planning (i.e. generation of a composite service from a user goal and a user profile). In its current version, DAML-S does not systematically address all non-functional aspects of services such as pricing, payment, temporal availability, and reputation. Instead, the language allows a developer to incorporate these aspects by sub-typing pre-defined service classes and properties. In that sense, our work can be seen as a foundation for refining DAML-S and other similar languages so as to include non-functional characteristics.

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 match-making 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. With the increasing ubiquity, complexity and dynamism of services it is clear that this research is essential for at least two purposes. Firstly, for designing languages that describe entries within catalogues of services, and queries over these catalogues. Secondly, for 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.

In the long term, we expect that this work will lead to an extensible service advertisement language. Extensibility is a key requirement, since it should accommodate domain-specific characteristics and ontologies. Another research avenue that we plan to pursue, is that of *service specialisation*. Service specialisation underlies any efforts of service classification, which in turn are essential for structuring any catalogue of service offers. Whilst existing classification schemes (see section 2) rely on purely functional aspects, 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 substitution, 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.
- Caelli, W., Longley, D., and Shain, M. (1991). *Information Security Handbook*. Macmillan Publishers.
- 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/wsd1.asp>.
- Commerce One Inc. (2000). XML Common Business Library (xCBL). <http://www.xcbl.org>.

- DAML Services Coalition (2001). Daml-s: Semantic markup for web services. In *Proc. of the International Semantic Web Working Symposium (SWWS)*, Stanford University CA, USA. <http://www.ksl.stanford.edu/projects/DAML>, accessed 5-Nov-01.
- 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.
- 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](http://ecommerce.ncsu.edu/business_models.html).
- 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>.
- 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/>.
- Zeithaml, V. and Bitner, M. (1996). *Services Marketing*. McGraw Hill, New York, NY USA.