



A Conceptual Architecture for Semantic Web Services (Extended version)

Chris Preist
Digital Media Systems Laboratory
HP Laboratories Bristol
HPL-2004-215
November 26, 2004*

semantic web, web
services,
architecture,
agents

This paper is an extended version of one presented at the International Semantic Web Conference, 2004. In it, we present an abstract conceptual architecture for semantic web services. We define requirements on the architecture by analyzing a set of case studies developed as part of the EU Semantic Web-enabled Web Services project. The architecture is developed as a refinement and extension of the W3C Web Services Architecture. We assess our architecture against the requirements, and provide an analysis of OWL-S.

* Internal Accession Date Only

♦ Shorter version HPL-2004-214, ISWC, International Semantic Web Conference, 8-11 November 2004, Hiroshima, Japan

Approved for External Publication

© Copyright Hewlett-Packard Company 2004

A Conceptual Architecture for Semantic Web Services (Extended Version)

Chris Preist

HP Labs, Filton Rd, Stoke Gifford, Bristol BS32 8QZ, UK
chris.preist@hp.com

Abstract. This paper is an extended version of one presented at ISWC 2004. In it, we present an abstract conceptual architecture for semantic web services. We define requirements on the architecture by analyzing a set of case studies developed as part of the EU Semantic Web-enabled Web Services project. The architecture is developed as a refinement and extension of the W3C Web Services Architecture. We assess our architecture against the requirements, and provide an analysis of OWL-S.

1. Motivation

In this paper, we present an abstract conceptual architecture for semantic web services. The motivation behind this is

- To provide a conceptual model to aid in the principled design and implementation of semantic web service applications;
- To allow the principled application of existing semantic web service ontologies such as OWL-S and WSMO;
- To contribute to the evolution of such languages, or the design of new ones, through the identification of appropriate features such a language needs;
- To motivate the development of appropriate infrastructure to support a variety of semantic web service applications;

The main impetus behind this work has been provided by the European Union Semantic Web-enabled Web Services project (SWWS). This project aims to develop appropriate Semantic Web Services infrastructure to support four case-studies, inspired by the design principles proposed in the Web Services Modeling Framework [2]. Although these four case studies provide the primary source of requirements on our architecture, to encourage generality, we also consider a variety of other applications of SWS described in the literature.

We take the W3C Web Services Architecture (WSA) [14] as an important input and source of inspiration. This provides a conceptual model and context for understanding web service technology and its relationship with other components required to bring about a global web services network. In particular, it adopts the view that the semantics of a service is provided by a contract between the service provider and

service requestor. Such a contract can be explicit or implicit, and can be agreed in a variety of ways. It is important to note that the contract represents an agreement about what a service does; it does not necessarily need to have any legal status.

The potential role of contracts in supporting web service interactions is becoming increasingly recognised. The importance of contracts in supporting interactions between agents in open systems, especially e-commerce, has long been recognised by the agent community (For example, [5]). Furthermore, proposals have been made to use contracts in semantic web services [3]. The adoption of contracts as representing semantics in the WSA is further evidence of this trend.

Though the WSA adopts the contractual view, the implications of this are not explored as fully as is required if the semantics are to be made explicit. Specifically, it does not make a clear separation between the semantics of a service provider agent and the semantics of a particular performance of a service. Furthermore, it does not provide a model to show the entities and activities involved in agreeing the ‘semantics’ of a particular service interaction (though it does informally discuss this process.) The conceptual architecture we present builds on the WSA to provide appropriate separation of concepts and additional constructs. We adopt the WSA as it stands as much as possible and extend it as necessary. However, at times we have found it necessary to make small changes to the WSA to allow the separation of distinct semantic concepts. Where we do this, we explain and justify these modifications.

The WSA presents four different views on the architecture, focused around different core concepts; message, service, resource and policy. Our focus is primarily on the service-oriented view. In particular, we do not consider policy, security or reliability issues and how these are affected by the addition of explicit semantics. To do so requires additional analysis beyond the scope of this paper.

The paper is structured as follows. Firstly, we discuss the semantics of the word ‘service’, and consider the different ways it can be used. Next, we briefly describe the SWWS case studies, and refer to additional applications discussed in the literature. We then present a set of requirements on a semantic web services architecture, derived from analysis of this set of applications. Next, we briefly summarise the service oriented model of the WSA, and present our conceptual architecture as an extension of this. We assess the architecture against our requirements, present an analysis of OWL-S and conclude.

2. What is a ‘Service’?

The word ‘service’ can be used in several different ways. If a clear conceptual architecture is to be developed for semantic web services, we believe that it is necessary to make these different uses explicit.

1. ‘Service’ as provision of value in some domain

In the economy nowadays, a great variety of different services are traded; cleaning of windows, provision of stock quotes, shipment of crates, showing of films, giving

messages, translation of documents, location of out-of-print books, provision of a broadband connection of a certain quality etc. In each case, the customer receives something of value to them in return for a payment to the supplier. The service the customer receives can be described in some language about the domain of interest. Such a description does not need to refer to the way in which the supplier and customer interact. This is what an economist or businessperson would refer to as a service.

On the internet, services can be provided which involve some financial transaction, and so correspond closely to the above perspective. However we can broaden this view of service without losing what is essential. A service can be provided without any payment, and the value the service provides need only be 'in the eyes of the requestor', not necessarily something that would be considered to have monetary value in the commercial world. For example, a service may be the provision of a specific piece of information from a database, the provision of a plan to achieve some goal, or the assignment of some domain task to a given software application. Hence, we define a service to be *the provision of something of value, in the context of some domain of application, by one party to another.*

We also need to distinguish between a particular provision of value (for example, location of a particular out-of-print book) from the general capability to provide (for example, the ability to locate out of print books). We refer to the former as a *concrete service*, and the latter as an abstract service. Hence an *abstract service* is defined as *the capacity to perform something of value, in the context of some domain of application.* This is close to the definition of service given in the WSA.

2. 'Service' as a software entity able to provide something of value

This usage is common in the computer science and IT community. Papers often speak of sending messages to services, receiving results from services, executing services etc. While this is clearly appropriate in the context of web services and service oriented architectures, it can cause some confusion when mixing usage of this definition with the definitions above. For that reason, we believe that it is more accurate to speak of such an entity as being (part of) a service provider agent (SPA). We also use the term 'web service' to refer to a specific application with a WSDL/SOAP interface.

3. 'Service' as a means of interacting online with a service provider

The most common example of this usage is 'negotiation service'; A negotiation service provides the means to interact (negotiate) with a given actor. This usage is clearly different from usage 1, in that negotiation in itself does not provide something of value. However, the outcome of the negotiation may do so (and hence may be a service in the first sense.) For this reason, we refer to this as the provision of a negotiation 'protocol' or 'choreography'.¹

All three of these aspects of service provision are essential; the 'value' in some domain, the interactions involved, and the implementation of the entity able to deliver what is needed. However, we believe that for conceptual reasons it is best to make clear distinctions between them.

¹ A 'negotiation service' can exist in the first sense; this is a service where you 'subcontract' some negotiation problem to a service provider, and they negotiate on your behalf.

3. Example Scenarios

We now present five example scenarios. Of these, scenarios B, C and E are from case studies in the Semantic Web-enabled Web Services project.

Example A: Currency Conversion quotation²

The SPA provides a currency quotation for the conversion of a given amount of one currency into another currency, at a given time. The interaction is simple; the requestor sends a message containing 3 arguments (an amount and two currency codes) and receives a response with three arguments (a new amount, the exchange rate used and a timestamp.)

Example B: Design of a SPA to proactively manage Virtual ISP problems

A virtual ISP provides internet bandwidth to customers. Occasionally, some fault may occur resulting in the actual level of service going below that agreed with the customer in the Service Level Agreement contract. When this happens, a monitoring system generates an alarm and an associated Quality of Service report. Management of such a problem requires coordination of several internal applications, including a workforce management system (to timetable an engineer to fix the problem) and a customer relationship management system (to proactively inform the customer of the problem, and update them as the problem is solved).

In our scenario, each of the applications to be coordinated has a set of web-service interfaces associated with them. These allow access of and interaction with the application in a variety of ways. In addition, the applications may proactively send update messages which need to be handled. The designer has the task of creating a process description which, when executed, will call and coordinate the various applications in a way which will solve the problem and keep the customer informed. This process description then forms the internal logic of an SPA for handling a QoS alarm report.

The designer carries out their task by preparing a rough specification of the process the SPA should implement. They then use a browser, searching the local service library only, to locate web services able to provide the functionality needed at different stages in the process. By inspecting the specification of the inputs and outputs of the web services, together with the definitions of the services they provide, the designer constructs a process. The process, when run, will result in a series of web service invocations and responses which implement the required behaviour. The designer then ‘wraps’ the SPA by providing a description of the service it provides, together with its associated message exchanges.

Example C: Discovery, selection and execution of a customer notification service

A service requestor, OAS, needs to send a message to one of their customers, Janet. It prepares a description of the service it requires in some ontology. This description states that a message ‘M’ is to be sent to Janet either by txt or email, and payment is to be made by credit card. Other parts of the description (e.g. price) are left undefined. OAS contacts a discovery service provider with this service request,

² By Terry Payne: <http://www.daml.ecs.soton.ac.uk/services/SotonCurrencyConverter.html>

and receives back a list of providers able to provide txt or email and accept credit cards, but no information as to whether they are able to contact Janet or what price they charge. OAS then accesses the SPA description of each one of these, and contacts them requesting a quote. It sends a description of the service it is interested in as before. The SPA returns the description of the service, with previously undefined parts now defined. Specifically, the price is defined to be whatever the service provider charges for this service. If it has no means to contact Janet, then it informs OAS it cannot perform the service. OAS then selects between the possible SPAs, and requests the service by sending the appropriate message to the SPA. The SPA then sends 'M' to Janet, bills the credit card and sends a confirmation to OAS.

Example D: 'Smart' book buying from several online book suppliers

A book purchasing agent contains information about which books a user wishes to purchase, together with preference information as to how much they are willing to pay, how they can pay, how important rapid delivery is to them, etc. We will assume in this example that all the books are CS textbooks. The agent contacts a discovery SPA, providing a description of the service it requires (again, possibly with some information unconstrained or partly constrained. For example, it may constrain the 'books' in the service to be CS textbooks, and may specify the delivery address.) The discovery SPA returns a list of providers which satisfy (as far as the discovery SPA can determine) the criteria provided. The purchasing agent then enters into a conversation with each of the possible book SPAs. It determines if the required books are available from each and at what price, places them in the baskets of all SPAs and begins all the checkout process. At this point, it will get a quote for the whole order together with various delivery options from each. It can then select an SPA and delivery option based on the various quotes and the preferences of the user. It confirms this order, and cancels the rest. The service provider then ships the books.

Alternatively, a more sophisticated algorithm for getting the best deal is possible: exploring the purchase of various subsets of books from different suppliers, choosing the set of such contracts which minimises the price, and checking out from that set of SPAs. Note that this is a simple form of service composition – combining a set of different services to meet the user's needs.

Example E: Provision of a logistics supply chain [1]

A company requires the transport of a crate from Bristol to Moscow. It already has long-term contracts in place for land transportation of crates from Bristol to Plymouth, and from St Petersburg to Moscow. However, its usual supplier of shipping services is for some reason unavailable and it needs to rapidly locate and agree a replacement.

A company representative interacts with their service requestor agent to prepare a description of the service they require. This will include the crate dimensions and weight, the route required (Portsmouth to St Petersburg), the earliest dispatch date, the latest arrival date, the insurance terms required and the B2B messaging standards it can use (RosettaNet). This description is then used as a request to a discovery SPA, and the requestor receives a list of service providers able to ship from Portsmouth to St Petersburg and use RosettaNet, but no information on dates, prices etc. The requestor then sends the description of the required service to each SPA, and receives

back zero, one or more possible contracts which satisfy the requirements. (i.e. the ship leaves after the earliest dispatch date and before the latest arrival date, etc.) The contracts will also specify additional information which was undefined in the description of the request – notably the price. The requesting agent selects one (or queries their user) and agrees to it; the associated SPA confirms this. The other SPAs are informed they will not be required.

Execution of the service is coordinated by a conversation. It starts when the crate is about to arrive in Portsmouth, with a RosettaNet Advanced Shipment Notification exchange, and ends with the sending of a Proof of Delivery and Invoice by the SPA when the crate arrives in St Petersburg. Full details of this conversation are beyond the scope of this paper.

This scenario requires a particularly rich form of ‘service composition’. Firstly, during service selection and agreement, it is important to ensure that the new logistics service can ‘slot in’ between the road transportation service to Portsmouth and the train service from St Petersburg. This means that the service descriptions contained in the three contracts must be such that one contract results in the delivery of the crate to a destination in time for the next service to collect it (but not too early that they refuse to accept delivery). Secondly, the three post-contractual conversations must be coordinated so that messages received during the execution of one service by the service requestor are processed and their contents used to generate messages transmitted during the execution of the next service. Hence, this case study requires both composition of service agreements, and choreography of post-contractual conversations during execution.

4. Requirements Analysis

Based on the above case studies (and others) we identify requirements on a conceptual architecture (CA) for semantic web-enabled web services³:

DISCOVERY

1. The architecture should define the functionality of service discovery, but should not constrain the way in which it is enacted.
2. Discovery should be possible either using an explicit description of a required capability (E) or an implicit description through 'service provider type'.
3. Discovery should identify a set of service providers which can potentially meet the requestor's need. (C,D,E)
4. It should be possible to do discovery through informal browsing by a user (B).
5. Discovery should be symmetric; the functionality should not restrict a provider to advertise and a requestor to actively discover.

³ The letters in brackets give the examples in section 2 which generate the requirement.

6. At the end of the discovery process, there may (D) or may not (A) be a need for a contract agreement conversation to define the exact service to be provided.

CONTRACT AGREEMENT

7. The initiator should know prior to contacting them whether a contract agreement conversation is necessary with the discovered agents to determine the details of the service.
8. A contract agreement conversation may involve negotiation of service parameters, selection of parameters from a set of possibilities (E), or a 'constructive' process such as the 'shopping trolley' metaphor (D).
9. A CA should enable negotiation to take place, but not place any requirements on either the structure of the contract agreement conversation required or on the negotiation strategies any participant may use. In particular, it should not require either the revealing of preference information or strategy adopted by either party.
10. At the end of a contract agreement conversation, there will be an explicit description of the agreed service to be provided, providing the 'semantics' of the agreed service. This may be contained in a contract. If there is no such conversation, then the outcome of the discovery process should provide something identifiable as an explicit representation of an agreed service.

SERVICE DELIVERY

11. A service may be delivered directly as an output of an interaction with a Service Provider Agent (A), immediately as the effect of an interaction with a SPA (C), or as a later event or sequence of events after the defining interaction (E).
12. A post-agreement conversation may be directly linked to the execution/delivery of an agreed service (E) or may simply monitor the progress of service delivery without directly affecting it (D).
13. A conversation about an agreed service may be a simple request-respond pair (A), or may be a more complex interaction(D).
14. Some conversations about an agreed service may be initiated by the requestor (A), some by the provider (E), and some by either party (C).
15. A conversation about an agreed service may involve an asynchronous exchange of messages (E).
16. The semantics of a post-agreement conversation should be conceptually different from the domain-level semantics of the agreed service itself. However, there should be a means of linking the conversation with its domain level effect (E).

SERVICE OUTSOURCING AND COMPOSITION

17. A service provider agent may in turn act as a service requestor agent to 'outsource' parts of the functionality it needs to deliver its service.

- 18. Some SPAs may achieve this outsourcing by executing processes of more simple elements created statically or dynamically using an internal planner (B); other SPAs may use very different reasoning methods (C).
- 19. It may be necessary to 'compose' agreed services (as defined by their contract) either to enable them connect in some way (E) or to provide a 'service bundle' that meets some need (D).
- 20. It may be necessary to 'compose'/coordinate conversations about services, to allow two services to interact appropriately (E).
- 21. Parts of the functionality required to support a conversation may themselves be treated as component services (D).

5. The Conceptual Architecture

In this section, we present a conceptual architecture which aims to meet the requirements listed in the previous section. The presentation is based on that used in the W3C-WSA.

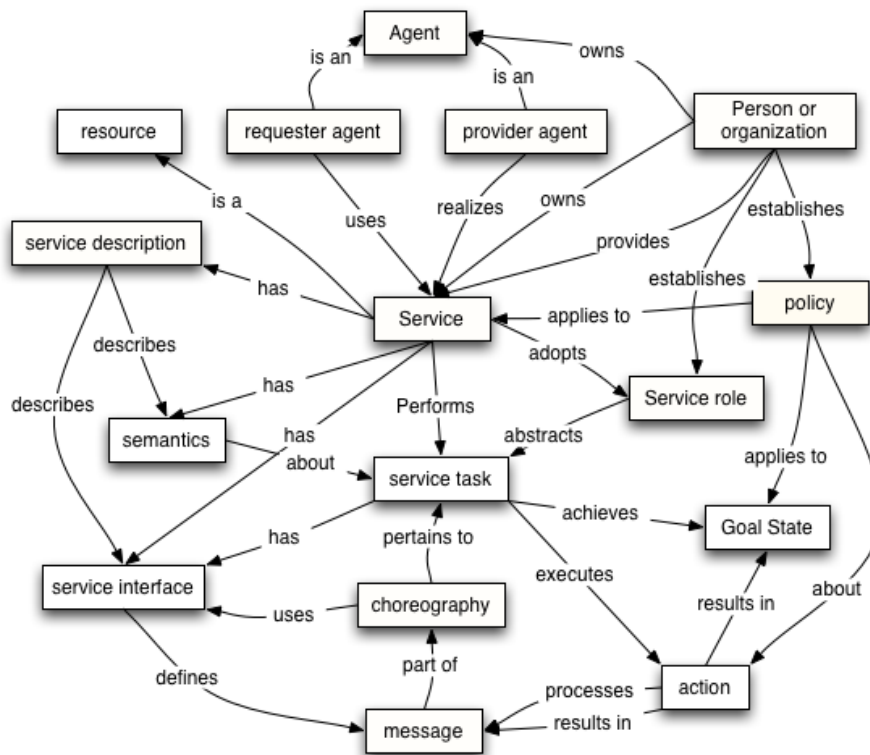


Fig. 1. The Service-oriented model of the W3C Web Service Architecture

Fig. 1 shows the service-oriented model of the WSA ([14]). We now give a brief overview of it. The concept of a service is core, which is defined as an abstract resource representing a capability to perform some coherent set of tasks. The performance of a given task involves some exchange of messages between the requestor and provider agents. A choreography defines a pattern of possible message exchanges (i.e. conversations). The WSA associates a choreography with a task's performance. The task may also have effects in the virtual or real world. Such effects may be publicly observable, but will often be private (such as a database update). The success or otherwise of a given task (from the perspective of one party) can be assessed by whether some goal state has been achieved. Such a goal state can be defined formally or informally.

As Sycara et al observe [10], tasks can be defined in three ways. They can be represented explicitly, using name labels with a well defined semantics in some ontology of tasks. They can be represented implicitly, using a language of preconditions and effects of a task. Or they can be defined using a combination of these two approaches. The WSA makes no commitment as to which approach is used. The primitive operation within tasks is referred to in the WSA as an action. However, the authors observe that 'the actions performed are largely out-of-scope... but the resulting messages are in scope.'

A service has a description, which specifies its interface and messages in a machine-processable way. This may also include a description of the service's semantics. The semantics specifies (formally, informally or implicitly) the intended effect of using the service: specifically, the tasks that constitute it. The WSA views this semantics as 'a contract between the requestor entity and provider entity concerning the effects and requirements pertaining to the use of a service.'

We adopt the majority of this model unchanged. The two changes we make are:

1. Distinction between abstract service, agreed service and concrete service

In the WSA, a service has a semantics which can be described, advertised and discovered by a service requestor. However, the service semantics is defined as the contract between the provider entity and requestor entity. This creates a problem: How can something which is defined as an agreement between the two entities be advertised prior to the two entities communicating? We believe this can be solved by making a distinction in the model between the *abstract service* a provider entity offers and the *agreed service* which appears within a contract, and assigning semantics to each. Additionally, (as outlined in section 2) we introduce the *concrete service* which is the performance of a specific service at a particular time. We give more complete definitions and examples of these concepts below. Similarly, we consider *abstract tasks*, *agreed tasks* and *concrete tasks*.

2. Distinction between service and service provider agent

The WSA specifies that a Service Provider Agent realises a service. The word 'realises' (and the explanation in WSA) implies that the provider agent is an embodiment of the service. We believe that this confuses a service as the capability to provide something of value with the entity or software implementation which has the capability. For this reason, we say that a provider agent *provides* an abstract or agreed service, and *performs* a concrete service.

We now present our conceptual architecture (CA). Firstly, we present definitions of the concepts we introduce. Then we present and explain the architecture. For clarity, we separate it into three diagrams, representing the entities and relationships involved in discovery, service definition and service execution. We omit the entities and relationships which appear in the WSA if they remain unchanged in our model and are not relevant to the discussion of service semantics. We adopt the notational conventions used by the WSA; in particular, we provide textual descriptions and diagrammatic representations of concepts and relationships between these concepts. Many of the relationships are descriptive, but ‘has’ is treated as definitorial. In the diagrams, boxes represent concepts and arrows represent relationships. The arrow points from the subject of the relationship to the object.

Firstly, we define the concept of a concrete service. This is the core ‘building block’ with which to define the semantics of more abstract descriptions of services.

Definition: A *Concrete Service* is an actual or possible performance of a set of tasks that represent a coherent functionality (and therefore deliver some value) within some domain of interest to its associated requestor and provider entities.

A concrete service:

- has a set of one or more concrete tasks
- has a set of messages
- has a service provider and service requestor
- has one or more post-agreement choreographies associated with it
- has a service provider interface and/or a service requestor interface which are used by the choreographies
- may have a service description
- may be a resource

A concrete service is effectively a (possible or actual) specific and detailed set of tasks carried out at a given time, and the messages exchanged in association with this. The messages are exchanged between a specific provider and requestor, according to certain choreographies and using certain interfaces.

An example of a concrete service would be: The service provider PayneConverter performs the task of quoting a euro exchange value for £10 for the service requestor cwp at 2pm on May 25, 2004. The message exchange consists of the requestor sending the message `exchangeRateRequest(10,uks,euro)` and the provider replying with `exchangeRateResponse(14.9599, 2004.05.25 14:00:29 GMT)`.

A service provider, however, does not offer a specific concrete service. Rather, they are able to provide many different possible concrete services. Similarly, a requestor is rarely after a specific concrete service; they are interested in many possible services which meet their needs. For that reason, we define the concept of abstract service:

Definition: An *Abstract Service* is some set of concrete services. A concrete service is said to be a *realization* of an abstract service if it appears in this set.

Definition: An *Abstract Service Description* is some machine-processable description D which has, as its model, an abstract service C . (We write $\text{Model}(D) = C$)

An Abstract Service description is some specification of a set of services. The format of this description is not specified as part of this architecture. There are several possibilities:

1. An extensional listing of concrete services.
2. A label representing the name of some abstract service in some ontology of services.
3. A template or formal description in some knowledge representation language.

Option one is the most conceptually straightforward, but in practice will be infeasible except in the simplest of cases. In option two, the semantics of the label is defined as the (possibly infinite) set of concrete services which correspond to that label within the ontology. This correspondence may be well-defined and explicit, with a service specification or contract template associated with the label. Often, however, the correspondence will be more informal. For example, a label ‘bookselling’ may have as its semantics the set containing any possible concrete services where books (and only books) are being sold.

The third option specifies the set of concrete services by explicitly and formally defining an abstract service using some knowledge representation (KR) language such as OWL or F-logic. This may involve explicit descriptions of abstract tasks, possibly including parameters. The set-theoretic semantics of the KR language then defines the set of concrete services to be those which formally satisfy the concept description. Such a representation may leave the choreographies to deliver the tasks undefined. In such a case, the abstract service will include all possible choreographies which can be associated with those tasks. Alternatively, the abstract service may place restrictions on the choreography (e.g. it may be constrained to be a RosettaNet PIP) or may define a specific required choreography.

A resource representing an abstract service will have an abstract service description associated with it. In addition, it:

- may have a contract agreement choreography associated with it
- may have a service provider, and an associated service provider interface
- may have a service requestor, and an associated service requestor interface

Whereas a concrete service is always associated with a specific requestor and provider, an abstract service may have one or both of these undefined. Similarly, an abstract task may be less completely defined than a concrete task is.

An example of an abstract service would be: The service provider *PayneConverter* can perform the task of quoting an exchange value for any amount, between any two currencies, at any time in the future. The message exchange will follow the choreography defined by its OWL-S process model. The abstract service is the set of all concrete services which meet this description. (Furthermore, a formalized version of this would be an example of an abstract service description.)

An abstract service, as we shall see below, can be used for advertising or requesting. However, it is not precise enough to define what a requestor and provider agree. For that we define the following;

Definition: An *Agreed Service* is an abstract service agreed between two parties. Hence it is associated with a specific service provider and service requestor, together with their associated interfaces. It also has specific post-agreement choreographies associated with it. Its definition is such that;

- (i) Any further decisions necessary to select a realization will occur as choices made during a post-agreement choreography.
- (ii) Any choice made by one party during the post-agreement choreography is such that the other party is indifferent to how that choice is made.

An example of an agreed service would be: The service provider PayneConverter agrees to perform the task of quoting a euro exchange value for any amount of UK sterling for the service requestor cwp any time before 12am on May 26, 2004. The message exchange will follow the choreography defined by its OWL-S process model.

PayneConverter is indifferent as to the exact value it converts and the time of conversion. Its choreography definition puts the decision of how much (via a parameter) and when (via choice of when to send the message) in the hands of cwp. Hence this is indeed a possible agreed service.

The agreed service forms the core of the service contract between two parties, which defines the semantics of their interaction:

Definition: A *Service Contract* is an agreement between a service provider and requestor that the provider will supply an agreed service to the requestor. The contract:

- has a service provider
- has a service requestor
- has an agreed service description
- has an association between the tasks involved in the agreed service description and choreographies and message interfaces which will be used to have conversations about the tasks.
- may have information about who is responsible and what compensation may occur in the case that a service is not delivered correctly. Compensation may involve remedial actions and/or payment of damages.
- may have additional legal terms and conditions.

The agreed service description will be (explicitly or implicitly) in terms of the tasks to be performed. It may include tasks to be performed by the requestor (e.g. payment transfer in the case of e-commerce applications.) The task/choreography association defines how the interaction between the parties can take place during delivery of the agreed service. This may be done by including the choreographies explicitly in the contract. This is appropriate where a choreography may vary and require agreement

between two parties. Alternatively, it may be done by referencing a choreography definition elsewhere. This is appropriate when the service provider agent offers a specific choreography and is not able or willing to alter it, or provides a choice between a small number of choreographies, or the parties use some standard choreography defined elsewhere.

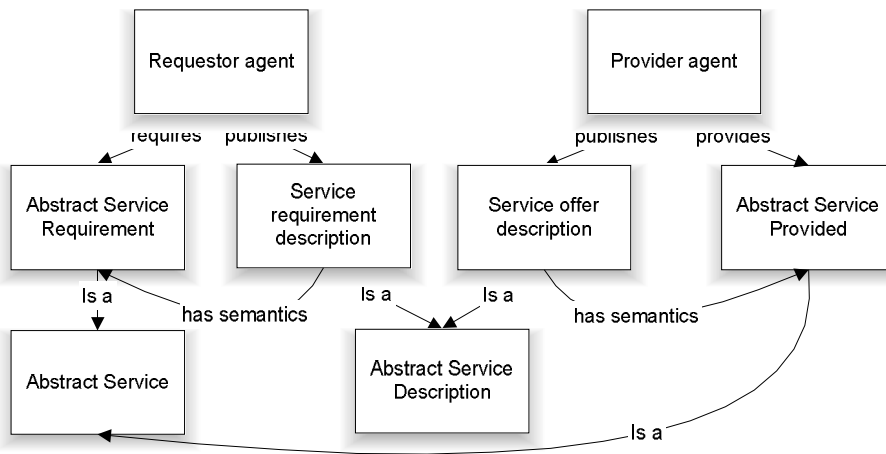


Fig. 2. Discovery Model of Conceptual Architecture

Fig. 2 illustrates the concepts involved in discovery. A requestor agent requires one of a class of abstract service; to this end, they publish a description of their requirements. Similarly, a provider agent is able to provide services within a certain class of abstract service. It publishes a description of the services it can provide.

Ideally, in each case, the description of the abstract service will correspond exactly to the actual abstract service which is required or offered. However, in practice, this may not be possible. For this reason, we define correctness and completeness of descriptions:

Definition: A description D of an abstract service C is said to be *correct* if any element in the model of D also appears in C . i.e. $Model(D) \subseteq C$.

Definition: A description D of an abstract service C is said to be *complete* if any element of C also appears in the model of D . i.e. $C \subseteq Model(D)$.

Often, it is possible to achieve completeness but not correctness in service descriptions. For example, a bookseller can state ‘I sell books’ in an advert, but cannot guarantee that every title is in stock. The role of the discovery process is to determine if a given service provider may be able to meet the needs of a given service requestor, and to inform one or both parties of this. As the WSA discusses, discovery can be achieved in several different ways, including the use of a discovery service. Note that we say that both the requestor and the provider publish a service

description. Such a publication may be as an entry into some discovery service, as a query to such a service, or as a local advertisement or request in a peer-to-peer system [2]. Following requirement 5, we make no commitment as to which party advertises and which party queries. In general, we would anticipate that providers would advertise and requestors would query the advertisements; however, there are exceptions such as when a requestor issues a public Request For Quotes describing the kind of service it requires.

Assuming that descriptions of service requirements and service offers are complete but not correct, then the minimal functionality required of the discovery system is to determine if the two abstract service descriptions intersect [11]. If this is the case, an agreement between the two parties may be possible. Paolucci et. al. [8] describe a more sophisticated matchmaker which uses a hierarchy of different definitions of matching to determine which matches are more likely to be fruitful.

Following discovery, two parties need to determine if they can agree on a service which the provider will supply to the requestor. This is done through the contract agreement phase. In some cases, this phase will be unnecessary and the results from the discovery phase will be enough. This is the case if the following criteria are satisfied;

- The published service offer description is correct (but not necessarily complete).
- The offer description provides a complete and explicit service description which includes post-agreement choreographies.
- Any aspect of the offer description which has variability (for example, a parameter which is not defined) is such that the provider is indifferent with regard to the choice.
- The choreography used to deliver the service is structured such that choices within this variability (such as a parameter instantiation) can be input through it.

If all four of these criteria hold⁴, then the contract agreement phase is unnecessary. The service offer description defines a set of concrete services, and the provider doesn't care which is selected. The delivery choreography allows the selection of a concrete service to be made by inputting parameter values chosen by the requestor. We call a service which has an abstract service description satisfying these criteria a **basic** service.

If the published abstract service description does not satisfy the criteria, but the service provider agent is able to give an abstract service description which does, then the contract agreement phase becomes a straightforward 'take-it-or-leave-it' process. The contract agreement choreography simply consists of the requestor accessing the description. It then goes on to execute the delivery choreography if it chooses to use this service provider, and does nothing otherwise.

⁴ This is phrased assuming the provider advertises and the requestor queries and selects. However, it is equally possible (though a lot rarer) that the requestor is indifferent, and the provider makes the selection using the requestor's delivery choreography.

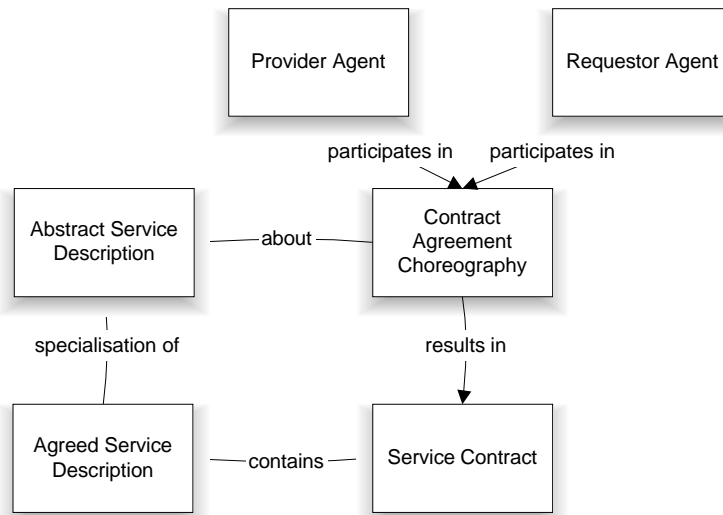


Fig. 3. Contract Agreement Model of Conceptual Architecture

In general, the contract agreement phase can be more complex than this. Fig. 3 illustrates the concepts it involves. At least one provider and at least one requestor agent participate in a choreography to agree a contract between them. The choreography involves discussing some abstract service, and may explicitly involve the manipulation of a description of that abstract service. If the choreography terminates successfully, it results in a service contract agreed between two parties; this contract contains a description of an agreed service which specifies the concrete service to be delivered. The agreed service will be a specialisation of the abstract service found during discovery. The decision making within the contract agreement phase has three aspects:

1. Partner selection

Often, one or both parties have a choice about which partner to select to deliver/receive a service. Specifically, it is often the case that a service requestor has a choice of several potential providers. Similarly, a service provider that has limited resources and is in demand may need to make a choice to supply some subset of requestors currently demanding its services.

2. Service Selection

One party may need to select and specify certain parameters of a service which weren't specified during discovery. Usually, this is the requestor making their requirements more specific (e.g. specifying exactly which books they wish to purchase from an online book-seller.)

3. Negotiation

Often, when a service specification is fully defined, other parameters of a contract such as price will be fixed. However, in some circumstances, they may be open to

negotiation. This means a previously proposed price (or other parameter) may be relaxed by one party in an effort to reach an agreement.

Any given contract agreement phase does not necessarily involve all three of these. Furthermore, the order in which these take place is not fixed; they may take place one after another in any order, or all three aspects can be interleaved with each other. In the agent mediated electronic commerce literature, a clear distinction is made between a negotiation protocol, defining the rules of negotiation, and the negotiation algorithm executed by a given agent [4]. Similarly, we distinguish here between a contract agreement choreography (in which two or more agents participate) and the algorithms within requestor and provider agents which they use to participate in such a choreography. The choreography is within the scope of the architecture, as it is inter-agent. The algorithms used by the agents, however, are private and so fall beyond the scope of the architecture.

Definition: A *contract agreement choreography* is a protocol for message interactions between at least one service requestor and at least one service provider which, on successful termination, will result in the agreement of a service contract between two of the parties.

A contract agreement choreography may explicitly involve service selection and/or negotiation. Service selection protocols can take different forms. Examples include:

1. A service provider may offer a list of all contracts they are willing to offer, and the requestor simply selects one.
2. In the common ‘shopping trolley’ conversation, a requestor iteratively builds a definition of the service they wish to receive, and are given a price quote when they are satisfied. They can accept the quote, quit, or make alterations to the required service.
3. ‘Pseudo negotiation’ may take place where partial descriptions of services (possibly using constraints on parameters) are exchanged between both parties in a series of proposals/counter-proposals. Slowly, by iterating, one party (usually the requestor) makes selections within the range and moves towards a fully instantiated contract. In this case, the provider has a set of fixed-price contracts (as example 1) but chooses to reveal these iteratively through a process which appears like negotiation. Sometimes this can be to make life easier for the requestor, as the provider co-operatively reveals contracts based on the revealed needs of the requestor. Sometimes, it can be in an attempt to guide the requestor to the provider’s preferred options.

Example negotiation protocols include:

1. Propose/counter-propose – Each party proposes a contract, and the other party either accepts it or proposes an alternative.
2. Constraint relaxation – Parties exchange a set of constraints on a contract template. Agreement occurs when the union of the two sets is satisfiable. Until this point is reached, parties iteratively relax their set [12].

3. Reverse Auction – A requestor sends out a call for quotes to fulfill a given service contract; providers respond with their cheapest quote, and iteratively place lower offers until no provider is willing to reduce any more. Agreement is made with the cheapest.

Many other possible approaches to negotiation exist ([4], [13]).

A choreography may also combine the two – for example, we could imagine a book selling service provider in which you use a shopping basket to choose your books, and then use propose/counterpropose to negotiate the final price. Usually, choreographies specifying these protocols will involve two parties (requestor and provider). However, either party may be involved in several simultaneous choreographies to allow them to make the optimal trading partner selection (Example D). Trading partner selection, therefore, is a consequence of the internal algorithm rather than the choreography. These parallel conversations may be hidden (as in example D,) or may be deliberately made public as in the reverse auction. In the latter case, the conversation protocol can be viewed as part of a multi-party choreography. The definition of the choreography would be made public to all participants.

The association of a contract agreement choreography to a given set of provider and requestor agents must take place following discovery. The most common way to do this is for the advertising agent to associate a (pointer to) a contract agreement choreography description with their advert, or to return a contract agreement choreography when they are first queried directly. The contract agreement choreography may take the results of the discovery phase as input, if the discovery phase revealed more information than just possible service providers. For example, the discovery phase may return a contract template for each service provider, and this can then be instantiated during the choreography. Alternatively, a provider agent may provide a new (and more detailed) abstract service description at the beginning of the choreography, or may allow a requestor to iteratively construct a service contract.

When a contract agreement choreography has terminated successfully, then service delivery can take place. This may commence immediately on agreement of the contract, or may be initiated later by a subsequent choreography. Fig. 4 illustrates the concepts involved in the service delivery phase, most of which have already been defined. A concrete service is supplied by the provider agent to the requestor agent. This concrete service must be a member of the agreed service (which, recall, is a set of concrete services.). The agreed service has an associated service contract which contains a description of the service agreed together with a mapping from the agreed tasks within the service to choreographies about these tasks. To distinguish these choreographies from the contract agreement choreographies, we refer to them as *post agreement choreographies*:

Definition: A post-agreement choreography is a protocol for a conversation between a service requestor and service provider about a contract which has been previously agreed, or one or more tasks within such a contract.

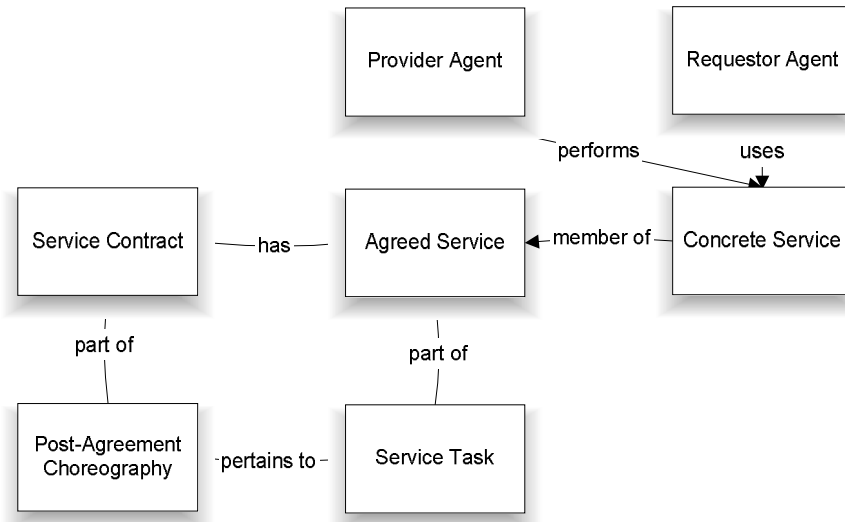


Fig. 4. Service Delivery Model of Conceptual Architecture

There are at least three different kinds of post-agreement choreography:

1. **Delivery choreography** - This governs conversations which directly initiate and/or control the execution of one or more tasks in a contract. For example, a request/response to an information service provider returning the required information, or a payment protocol transferring money from one party to the other.
2. **Monitoring choreography** - This does not directly control the execution of a task, but allows the requestor to monitor the task's progress. In example E, the logistics provider allows the requestor to send status request messages, and returns information about the package's current location and estimated arrival time.
3. **Cancellation/renegotiation choreography** - This allows the cancellation or alteration of a service within a contract in certain circumstances. For example, a book seller such as Amazon may provide the option to cancel or alter an order after it has been submitted, providing the order has not entered the dispatch process. This may involve compensation.

This completes the Semantic Web-enabled Web Services conceptual model (CA). We now assess the architecture against the requirements of section 4.

6. Assessment of the Conceptual Architecture

Requirements 1, 2, 3 and 5 are straightforwardly satisfied by the structure of the conceptual architecture applied to the discovery phase as discussed above. Requirement 4, that discovery can be performed by informal browsing as well as

matching, is not satisfied immediately by the architecture. However, a browser can easily be implemented within the conceptual framework described above; a set of service offer descriptions can be filtered, ordered and presented to a user according to certain parameters specified by a user.

To meet requirement 6, the architecture allows but does not require a contract agreement choreography between the parties. Furthermore, we characterise services which do not need such a choreography, and define them as basic services. While this characterisation is not technically part of the architecture, a discovery service could use it to determine if a contract agreement choreography will be necessary, satisfying requirement 7. Requirements 8, 9 and 10 are straightforwardly satisfied by the structure of the conceptual architecture as applied to the contract agreement phase as discussed above. In the case of basic services, an explicit contract is provided by the service offer description which is offered as take-it-or-leave-it and so needs no further refinement.

The architecture makes a clear separation between the definition of a service, interactions defining the service, and interactions involved in the delivery of the service. In particular, the delivery of a service can take place independently of its definition, satisfying requirement 11 and part of 16. The requirement that there be a means of linking a conversation with its effect is enabled but not satisfied by the architecture; use of a language which allows this is necessary. The other requirements on service delivery choreographies (12, 13, 14, 15) are straightforwardly satisfied by the structures defined above.

The conceptual architecture, like the WSA, allows ‘outsourcing’ of functionality (17) as agents can act simultaneously both as service providers and service requestors, possibly of several different services. We place no restrictions on the internal reasoning used to determine how this is done. One option is to use a planner to determine which service requests to make (18).

By separating out the definition of an agreed service from the choreographies to define and deliver the service, the architecture allows different kinds of ‘composition’ to take place. In particular, it enables parties to reason over possible service definitions at the domain level during the contract agreement phase to ensure coordination between the different services agreed (19). It also enables parties to reason about, and coordinate, choreographies associated with the agreement and deployment of different services (20). However, both of these are merely enabled by the architecture; reasoning strategies to do these must be embodied within specific requestor/provider agents.

A third form of what could loosely be called ‘composition’ (21) is the application of some service to enable a choreography between provider and requestor of some other service. For example, a ‘shopping trolley’ choreography to define a book purchase may be enabled by component services, with functionality such as ‘book search’ or ‘put in trolley’. The domain of application of these component services is ‘shopping trolley choreography’ as opposed to ‘book delivery’, the domain of the ‘higher level’ service being defined. However, within their domain, they can be considered as services with associated contractual definitions, and can be located and reasoned over by an agent which brings about the choreography between requestor and provider of the book delivery service. Such an arrangement, as opposed to fixed

choreography definitions, would enable more flexible interactions between the two parties. It is conceivable to imagine a hierarchy of services nested in this way; though in the end it would bottom-out to fixed choreographies or basic services.

Similarly, mediation services will play a part in enabling other services to be defined and delivered. Such services are a core aspect of WSMF [2], which identifies many different roles mediation services can perform. A full analysis of how these different forms of mediation are embodied in the conceptual architecture requires further work and is beyond the scope of this paper. However, we can identify the following requirements for mediation in our architecture;

- The ontology language used for a service offer description may be different from that used for a service request description.
- The message used in a choreography by one agent may have a different format, but similar meaning, to that used by the other agent.
- The two agents may require different choreographies.

In each case, an agent could provide a mediation service to support the required interaction.

7. Analysis of OWL-S with respect to the Conceptual Architecture

OWL-S ([7],[9]) (formerly DAML-S [6]) is an ontology for use in providing semantic markup for web services, developed by a coalition of researchers as part of the DARPA Agent Markup Language program. In this section, we compare the conceptual model implicit in OWL-S with our conceptual architecture.

The OWL-S ontology defines a service in terms of three top-level classes; the profile, the service model and the grounding. The profile characterises what a service requires and provides, and is used primarily in the discovery phase. The service model describes how a service works, primarily in terms of a process model. This defines a service's functionality as the execution of a (potentially complex) process of component web services. The grounding defines how a component web service can be interacted with, primarily through WSDL/SOAP.

These three classes roughly correspond to our requirement for a service definition at the domain level, a description of how to interact with a service provider, and an implementation of that interaction. However, there are some differences which mean that OWL-S would need to be adapted and augmented if it were applied within our conceptual architecture.

In OWL-S, the profile is used almost exclusively as an advertisement/request. Hence it focuses on the description and categorization of an abstract service, but does not consider what is necessary to refine it to define an agreed service and to use this as a specification of what interactions accomplish. To do this, the service model (either within or outside the profile) would need to be augmented to provide complex structured descriptions of abstract services, and the ability to constrain and instantiate service parameters.

The profile does not make a clear distinction between what a service does and the way in which one interacts with a service provider to receive the service. In particular, it is not clear that inputs/outputs should appear in the domain-level description of an agreed service. (For example, the domain level agreed service description may state that ‘provider pays £10 by credit card’; the associated choreography will require a credit card number as its input. Hence the domain level definition and the choreography to deliver it are kept separate.)

The process model in OWL-S appears to serve two separate functions. Firstly, it can be used to define the exchange of messages with a service provider about a service; secondly, it can be used to define how a service provider implements the functionality of a service as a process of component web services. Both of these are valid uses, but the conceptual architecture requires that they be separated out.

The required exchange of messages about a service is defined by one or more post agreement choreographies in our conceptual architecture. The process model could be used to represent this. Note that the architecture can have several choreographies associated with a service, each performing different functions. Hence a given service may need several process models, and links between the process models and the domain-level service description to define what each does.

The definition of a service’s functionality in terms of a process over component web services is a particular method of delivering a service. The CA requires that such a method should not be proscribed and that a provider using such a method should not be required to make public their method. For this reason, using the process model in this way would not be compliant. A separate structure defining the internal process should be used, and the publication of it would be optional. (Such a separation is being adopted in the nascent WSMO [15])

8. Conclusions and Future Work

We have presented a conceptual architecture developed out of a requirements analysis of a set of case studies in the SWWS project and example applications from the literature. The architecture makes a clear separation between the domain-level definition of a service to be provided, the message exchanges which bring about the service, and the internal logic which provides the appropriate functionality. Inevitably, the architecture is at an abstract level; however, it does shed light on the required characteristics of service ontologies and infrastructure services to realise it. We are currently working on an extension of the CA to include sub-communities of service agents which share a problem-solving method (in particular, planning). We also hope to expand it to cover mediation more effectively.

The conceptual architecture is being used to define, design and implement a general technical architecture for Semantic Web Services able to demonstrate the four case studies within the SWWS program. It is also being used as the underlying conceptual model in the development of WSMO-full [15] and has been submitted for consideration by the architecture subcommittee of the Semantic Web Services Initiative.

Acknowledgements

The following people have contributed to this work through discussion and/or feedback: The HP Semantic Web services team; the EU Semantic Web-enabled Web Services team, Dieter Fensel, Mark Burstein, Ian Dickinson, Terry Payne, Bijan Parsia and the anonymous ISWC reviewers.

References

1. Esplugas-Cuadrado, J., Preist, C. and Williams, S.: Integration of B2B Logistics using Semantic Web Services. To appear, Proc. Artificial Intelligence: Methodology, Systems, and Applications, 11th International Conference, (2004)
2. Fensel, D. and Bussler, C.: The Web Service Modeling Framework WSMF. *Electronic Commerce: Research and Applications*, 1 (2002) 113-117
3. Grosz, B. and Poon, T.: SweetDeal: Representing Agent Contracts With Exceptions using Semantic Web Rules, Ontologies and Process Descriptions. To appear, *International Journal of Electronic Commerce* (2004)
4. He, M., Jennings, N.R. and Leung, H.: On Agent Mediated Electronic Commerce. *IEEE Transactions on Knowledge and Data Engineering* 15(4) (2003) 985-1003
5. Jennings, N.R., Faratin, P., Johnson, M.J., O'Brien, P. and Wiegand, M.E.: Using Intelligent Agents to Manage Business Processes. *Proceedings of the First Int. Conference on the Practical Application of Intelligent Agents and Multi-Agent Technology* (1996) 345-360
6. McIlraith, S. and Martin, D.: Bringing Semantics to Web Services. *IEEE Intelligent Systems*, 18(1) (2003) 90-93
7. OWL-S 1.0 Release. <http://www.daml.org/services/owl-s/1.0/>
8. Paolucci, M., Kawamura, T., Payne, T.R. and Sycara, K.: Semantic Matching of Web Service Capabilities. *Proc. International Semantic Web Conference* (2002) 333-347
9. Paolucci, M. and Sycara, K.: Autonomous Semantic Web Services. *IEEE Internet Computing*, (September 2003) 34-41
10. Sycara, K., Paolucci, M., Ankolekar, A. and Srinivasan, N.: Automated Discovery, Interaction and Composition of Web Services. *Journal of Web Semantics* 1(1), Elsevier (2003)
11. Trastour, D., Bartolini, C. and Gonzalez-Castillo, J.: A Semantic Web Approach to Service Description for Matchmaking of Services. In *Proceedings of the Semantic Web Working Symposium*, Stanford, CA, USA, July 30 - August 1, 2001
12. Trastour, D., Bartolini, C. and Preist, C.: Semantic Web Support for the Business-to-Business E-commerce pre-contractual lifecycle. *Computer Networks* 42(5) (August 2003) 661-673
13. Vulkan, N.: *The Economics of E-Commerce*. Princetown University Press, Princetown, New Jersey (2003)
14. W3C. Web Services Architecture. W3C Working Group Note, 11 February 2004. <http://www.w3.org/TR/ws-arch/>
15. Web Services Modelling Ontology: <http://www.wsmo.org/>