Annotation, Composition and Invocation of Semantic Web Services

Sudhir Agarwal, Siegfried Handschuh, Steffen Staab
Institute of Applied Informatics and
Formal Description Methods (AIFB),
University of Karlsruhe (TH),
D-76128 Karlsruhe, Germany.
{sag,sha,sst}@aifb.uni-karlsruhe.de
http://www.aifb.uni-karlsruhe.de/WBS

Abstract

The way that web services are currently being developed places them beside rather than within the existing World Wide Web. In this paper we present an approach that combines the strength of the World Wide Web, viz. interlinked HTML pages for presentation and human consumption, with the strength of semantic web services, viz. support for semi-automatic composition and invocation of web services that have semantically heterogeneous descriptions. The objective we aim at eventually is that a human user e.g. a consultant or an administrator can seamlessly browse the existing World Wide Web and the emerging web services and that he can easily compose and invoke Web services on the fly.

This paper presents our framework, OntoMat-Service, which trades off between having a reasonably easy to use interface for web services and the complexity of web service workflows. It is not our objective that everybody can produce arbitrarily complex workflows of web services with our tool, the OntoMat-Service-Browser. However, OntoMat-Service aims at a service web, where simple service flows are easily possible — even for the persons with not much technical background, while still allowing for difficult flows for the expert engineer.

1 Introduction

The Stencil Group defines web services as: loosely coupled, reusable software components that semantically encapsulate discrete functionality and are distributed and programmatically accessible over standard internet protocols. Though this definition captures the broad understanding of what web services are, it raises the question, what web services have to do with the web. Even if HTTP is used as a communication protocol and XML/SOAP to carry some syntax this appears to be a rather random decision than a deeply meaningful design.

We believe that it makes sense to actually integrate the strengths of the conventional World Wide Web, viz. lightweight access to information in a highly-distributed setting, with the strengths of web services, viz. execution of functionality by lightweight protocols in a highly-distributed setting. To seamlessly inte-

grate the two aspects we envision a *service web* that uses XHTML/XML/RDF to transport information and a web service framework to invoke operations and a framework, OntoMat-Service, to bind the two aspects together. OntoMat-Service offers an infrastructure, OntoMat-Service-Browser, that allows

- for seamlessly browsing conventional web pages, including XHTML advertisements for web services;
- for direct, manual invocation of an advertised web service as a one-off use of the service;
- for tying web service advertisements to each other when browsing them;
- for tying web service advertisements to one's own conceptualization of the web space when browsing them; and
- for invoking such aggregated web services.

For these objectives, we build on existing technologies like RDF [9], ontologies [1] or WSDL [22]. To integrate the web and web services into the service web, we make specific use of a new type of *semantic annotation* [5], namely *deep annotation* [6].

The paper proceeds as follows. We first describe a simple use case for OntoMat-Service (cf. section 2), including a detailed WSDL description of a web service used for the running example. In section 3, we describe the process that allows to turn web services into a service web and that lets a user browsing the web with OntoMat-Service-Browser exploit the very same tool to aggregate and invoke web services. The first step of this process, i.e. advertising web services in a form that combines presentation for human and machine agent consumption, is sketched in section 4. The second step of this process, i.e. using browsing and semantic deep annotation to tie together conceptual descriptions, is described in section 5. The third step comprises the generation of simple web service flows and is described in section 6. The fourth and final step described in section 7 deals with the invocation of web service flows. Before we conclude, we overview some related work.

2 Use Case

A typical use case supported by OntoMat-Service is the following (adapted from a larger scenario in [11]): Employees in an enterprise often need a new laptop. To make the laptop purchasing process easier for the employees, an administrator having technical knowledge defines a process for the employees of the enterprise. In order to buy a laptop it is desired to first collect offers from various laptop vendors based on the characteristics of the desired laptop like processor speed, disk size, etc. Further, it should be possible to close an insurance contract for a newly bought laptop. For this purpose insurance terms from a third party have to be collected. Once the most reasonable laptop and the best insurance contract terms are determined, the employee purchases the laptop and closes the service contract.

In our scenario, we assume a laptop vendor and an insurer offering web services with two operations each, i.e. `getLaptopOffer` / `buyLaptop` and `getInsur`

anceTerms / closeServiceContract, respectively. The sequence of operations that must be executed by the customer is depicted in Figure 1.

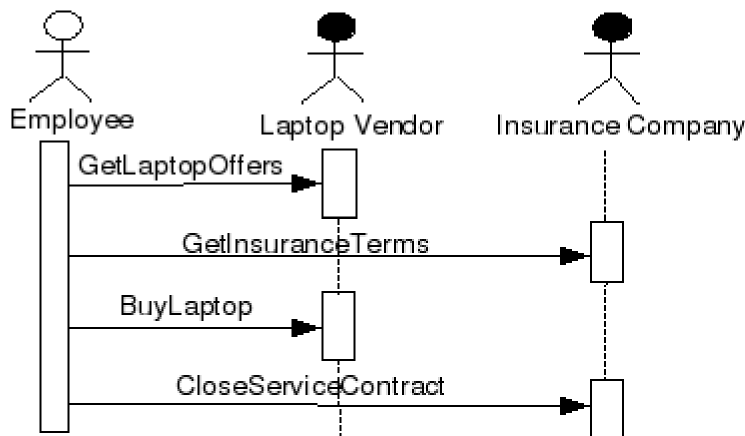

Figure 1: Sequence Diagram for the Use Case

The laptop vendor and the insurer being web service providers describe their web services with WSDL documents. In Figure 2, we show how a conventional WSDL document of the laptop vendor located at http://laptop-vendor.de/laptop.wsdl might look like.[1]

The WSDL document describes:

- *Data type definitions* in the XML element types. They are only sketched in figure 2 as they correspond to the laptop vendor's ontology depicted in N3[2] in figure 4. Thereby, we assume the definitions given in Figure 3. In our running example, the WSDL document of the laptop vendor, we describe the class Laptop.

- *Messages* that a service sends and/or receives and that constitute the web service operations in the XML element portType. For instance, our running example specifies 'getOffersRequest' that a potential customer would send to the laptop vendor to solicit an offer. getOffersRequest must be provided with two arguments, namely processor speed and disk size. It returns a set of laptop offers with properties such as specified in the vendor ontology (cf. WSDL document in Figure 2 and vendor ontology in Figure 4).

WSDL provides a naming convention for URIs such that each conceptual element (e.g., types, portType, etc.) of a WSDL document can be uniquely referenced. Such a URI consists of a targetNamespace pointing to the location of the

[1] The single ideosyncrasy we have here is that the WSDL document employs RDFS in order to describe the data structures instead of the more common XML schema — though actually WSDL does not require XML Schema and it allows RDFS.

[2] Notation 3 or N3 is basically equivalent to RDF in its XML syntax, but more compact. Cf. http://www.w3.org/DesignIssues/Notation3

```
<?xml version="1.0" encoding="UTF-8"?> <definitions
name="LaptopService"
    targetNamespace="http://laptop.wsdl/laptop/"
    <types>
       <rdf:RDF>
          <rdfs:Class rdf:ID="Laptop">
             <rdfs:label>Laptop</rdfs:label>
          </rdfs:Class>
          <rdf:Property rdf:ID="diskSpace">
             <rdfs:label>diskSpace</rdfs:label>
             <rdfs:range rdf:resource="&rdfs;Literal"/>
             <rdfs:domain rdf:resource="#Laptop"/>
          </rdf:Property>
           ...
          <rdf:Property rdf:ID="price">
             <rdfs:label>price</rdfs:label>
             <rdfs:range rdf:resource="&rdfs;Literal"/>
             <rdfs:domain rdf:resource="#Laptop"/>
          </rdf:Property>
           ...
       </rdf:RDF>
    </types>
    <message name="getOffersRequest">
        <part name="processorSpeed" type="processorSpeed"/>
        <part name="diskSpace" type="diskSpace"/>
    </message>
    <message name="getOffersResponse">
        <part name="laptopOffers" type="laptops"/>
    </message>
        ...
    <portType name="LaptopService">
        <operation name="getLaptopOffers" parameterOrder="processorSpeed diskSpace">
            <input message="tns:getOffersRequest" name="getOffersRequest"/>
            <output message="tns:getOffersResponse" name="getOffersResponse"/>
        </operation>
        ...
    </portType>
</definitions>
```

Figure 2: Web Service Description of Laptop Vendor

```
@prefix rdfs: <http://www.w3.org/rdf-schema#>. @prefix : <#>.
@prefix a rdf:Type.
```

Figure 3: N3 shortcuts

WSDL document and to element names of the WSDL document. For example, the URI `http://laptop.wsdl/laptop/#part(getOffersRequest/diskSpace)` refers to the second part (`diskSpace`) of the message `getOffersRequest` of the WSDL document in Figure 2 (cf. [22] for further specifications).

The web service description of the insurer looks similarly. We here only mention that the insurer provides the operations `getInsuranceTerms` and `closeServiceContract`. `getInsuranceTerms` requires a description of Laptop (according to the insurer's ontology in Figure 5) and a timePeriod, for which the contract is supposed to run. `getInsuranceTerms` returns a set of insurance terms available.

In the remainder of the paper, we assume that the customer has the plan depicted in Figure 1. However, in our running example, we will mostly focus on the first two steps to illustrate our framework.

```
:Laptop a rdfs:Class.
:price rdfs:domain :Laptop; rdfs:range :rdfs:Literal.
:diskSpace rdfs:domain :Laptop; rdfs:range :rdfs:Literal.
:processorSpeed rdfs:domain :Laptop; rdfs:range :rdfs:Literal.
:laptopID rdfs:domain :Laptop; rdfs:range :rdfs:Literal.

:Offer a rdfs:Class.
:laptops rdfs:domain :Offer; rdfs:range :Laptop.

:Sale a rdfs:Class.
:laptop rdfs:domain :Sale; rdfs:range :Laptop.
:creditCardNumber rdfs:domain :Sale; rdfs:range :Literal.
:customerReceipt rdfs:domain :Sale; rdfs:range :Literal.
```

Figure 4: Ontology of the laptop vendor

```
:Laptop a rdfs:Class.
:id rdfs:domain :Laptop; rdfs:range :Literal.

:ContractTerms a rdfs:Class.
:laptop rdfs:domain :ContractTerms; rdfs:range :Laptop.
:timePeriod rdfs:domain :ContractTerms; rdfs:range :Literal.
:price rdfs:domain :ContractTerms; rdfs:range :Literal.
```

Figure 5: Ontology of the insurance company

3 Overview of the Complete Process of OntoMat-Service

Figure 6 shows the complete process of our framework, OntoMat-Service. First, the figure consists of process steps, which are illustrated by a circle representing the step and a person icon representing the logical role of the person who executes the step, viz. service provider, annotating Service Web browser and a user invoking a Web Service. The two latter roles typically coincide. Second, the figure comprises information that is used by a person or by OntoMat-Service-Browser in a process step.

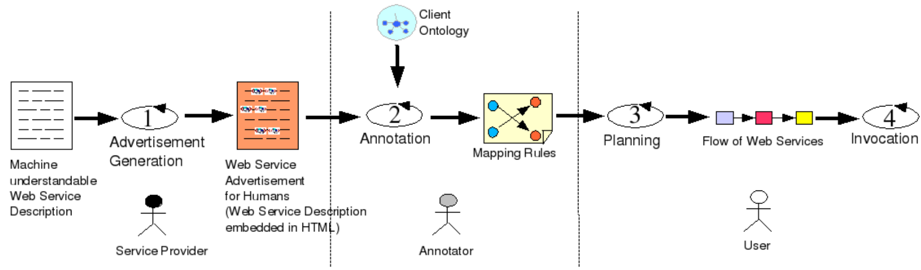

Figure 6: The Complete Process of OntoMat-Service

The four main steps run as follows:

Init: OntoMat-Service starts with a common WSDL web service description by the service provider (e.g., Figure 2). Obviously, the WSDL document is primarily intended for use by a machine agent or a software engineer who

has experience with web services. It is not adequate for presenting it to a user who is 'only' expert in a domain.

Web Service Presentation (Step 1): In the first step, the web service provider makes the web service presentation readable as a nicely formatted (X)HTML document — possibly including advertisements, cross-links to other HTML pages or services, or other items that make the web page attractive to the potential customer (cf. Section 4 for details).

Thereby, it is important that the understandable, but informal description of the web service is implicitly annotated to relate the textual descriptions to their corresponding semantic descriptions in their WSDL document.

Step 1 is a manual step that may be supported by tools such as *WSDL Documentation Generator* from *http://www.xmlspy.com*. However, we would not assume that tools like *WSDL Documentation Generator* would be sufficient to generate an amenable presentation, as they still produce rather rigid and technically oriented descriptions.

Result 1: Human understandable web page that advertises the web service and embeds/refers to machine understandable web service descriptions (WSDL + ontology).

Deep Annotation (Step 2): At a client side, a potential user of the web service browses the web page. OntoMat-Service-Browser shows the web page like a conventional browser. In addition, OntoMat-Service-Browser highlights human-understandable items (e.g. text phrases) that associate an underlying machine-understandable semantics.

The logical role of the user here is one of an annotator/browser. He can decide to just view the page and proceed directly to step 4 (described below). Alternatively, he can decide to map some of the terminology used in the web page of the web service to his own terminology (or to the terminology of someone else).

For the latter purpose, he loads an ontology into OntoMat-Service-Browser (if it is not already pre-loaded). Then he aligns terminology mentioned in the web page by drag'n'drop-ping it onto the ontology loaded into OntoMat-Service-Browser. OntoMat-Service-Browser generates mapping rules from these annotations that bridge between the ontology of the service provider and the ontology loaded into OntoMat-Service-Browser (cf. Section 5 for details).

Typically, the user will map to more than one web service, i.e. often he will map to different ontologies.

Result 2: Sets of mapping rules between web service ontologies and pre-loaded ontology.

Web Service Planning (Step 3): At the client side, a user might view the web services as well as their annotations that yield mapping rules. The third logical role here is one of a service planner and invocator (this logical role is shared between the third and fourth step). For this purpose, the user decides to select

- a set of web service operations he wants to use and
- a set of mapping rules he wants to use.

The reader may note that very frequently the roles of an annotator/browser and a service invocator will just coincide. Hence, the two selections just mentioned will take place implicitly — just by the web service pages he has browsed and the annotations that the service invocator has performed in step 2 of the OntoMat-Service process.

Once the two selections have been performed im- or explicitly, a module for web service planning will compute logically possible web service flows. For this objective, web service planning may employ a rich set of knowledge: goals, pre-conditions of web services, post-conditions of web services, previous similar cases, etc. In the current version of OntoMat-Service we just exploit the pre- and post-conditions derived from mapping one web service output to another web service input via the customer ontology. The web service description in the associated WSDL document describes what types are required for the input of a web service and what types appear in the output of a web service. Since data that wanders from one web service to the next can only proceed if types are compatible, OntoMat-Service-Browser can compute a restricted set of possible web service flows (cf. Section 6).

Though in general this model may be too weak to compute complex flows it is quite sufficient and straightforward to use with a small number of selected and semantically aligned web services — such as an end user or prototype builder will use.

Result 3: Sets of possible web service flows.

Web Service Invocation (Step 4): The final user, i.e. the invocator, can select one such flow from the list or modify any, if none of them fits his needs. Obviously, he can always create a new flow on his own. Once the user has a flow that fulfills his current needs, he invocates the flow (cf. Section 7). During the execution, the transformation of the data of one ontology to another will happen automatically via the mapping rules. The user achieves his goal at the completion of the invocation of the web service flow.

4 Semantic Web Page Markup for Web Services

In this section we show how a web service provider can semantically annotate the web pages describing his web services. Such a combined presentation allows for improved ways to find the web services (e.g., by a combined syntactic/semantic search engine) and it enables a user to understand the functionality of a web service and define mapping rules between the ontology used in the web service description and the client's ontology.

The basic idea is that a conventional HTML page about the web service and web service parameters is extended by URIs referring to conceptual elements of the corresponding WSDL documents. To carry these two pieces of information, we use `wsdlLocation` and `elementURI` inside the `span` tags. In Figure 7, we

```
<html><head><title>Laptop Vendor Service</title></head>
<body><h1 align="center">Laptop Vendor Service</h1>
<p><h2>getLaptopOffers</h2>
This service delivers the top offers of the laptops available in
the city. We have the largest archive of the laptop offers for the
city. So, the possibility that you find your desired laptop at a
reasonable price is very high. Just try it and get convinced from
our great offers. <ul>
<li> <span wsdlLocation="http://laptop-vendor.de/laptop.wsdl"
elementURI="http://laptop.wsdl/laptop/#part(
getOffersRequest/processorSpeed)">
<b>Processor speed</b> </span> Specifies the speed of the
processor. Please use only the units "MHz" and "GHz". For example,
"2GHz", "1.4GHz" and "1600MHz" are valid whereas "1800" or
"170000KHz" are invalid. </li>
<li> <span wsdlLocation="http://laptop-vendor.de/laptop.wsdl"
elementURI="http://laptop.wsdl/laptop/#part(
getOffersRequest/diskSpace)">
<b>Disk space</b> </span> Specifies the disk space. Please use
only the units "GB" and "MB". For example, "20GB", "30.5GB" are
valid whereas "40" or "25000KB" are invalid. </li>
<li> <span wsdlLocation="http://laptop-vendor.de/laptop.wsdl"
elementURI="http://laptop-vendor.wsdl/laptop/#part(
getOffersResponse/laptopOffers)">
<b>Top Offers</b> </span> This is the list of the most reasonable
offers available in the city that fulfill your requirements.
</li>
</ul></p>              ...
</body></html>
```

Figure 7: Web Service Description as HTML Page

show how such a web service advertisement(HTML page) for the laptop vendor service might look like.

When such an HTML page is opened in OntoMat-Service-Browser, the `span` tags are interpreted and elements between `<span>` and `</span>` are highlighted to support the annotation step described in the next section.

5 Browsing and Deep Annotation

In this section, we describe the second main step of the OntoMat-Service process. This step consists of browsing web pages about web services with OntoMat-Service-Browser. Thereby, the user may annotate [6] these web pages generating mapping rules between a client ontology and the ontologies referred to in the WSDL documents as a 'side effect' of annotation. We call this action 'deep-annotation' as its purpose is not to provide semantic annotation about the surface of what is being annotated, this would be the web page, but about the semantic structures in the background, i.e. the WSDL elements.[3]

Thus, this step is about web service discovery by browsing and using information retrieval engines like Google as well as about reconciling semantic heterogeneity between different web services, such as described in the WSDL documents and the web service ontologies they embed or refer to.

[3] [6] goes into detail for using deep annotation as the basis of database integration.

5.1 Service Browsing

With OntoMat-Service-Browser the user can browse the service web, i.e. he can browse the web pages of web service advertisements and OntoMat-Service-Browser highlights semantic annotations added by the web service provider. OntoMat-Service-Browser indicates semantically-annotated web service elements, e.g. input parameters, by graphical icons on the web page. Thus, the user may easily identify relevant terminology that needs to be aligned with his own ontology.

As an alternative to deep annotation, the ontology browser in OntoMat-Service-Browser may also visualize the underlying service ontology. OntoMat-Service-Browser is able to interpret the description of web service operations and provide a corresponding form interface (cf. figure 17). The user may then directly proceed to web service invocation (Section 7) and invoke a concrete web service operation with data he provides via this generic form interface.

5.2 Deep Annotation

The user selects an ontology to be used for annotation and loads it into OntoMat-Service-Browser. The user annotates the web service by drag'n'dropping highlighted items from the web page into the ontology browser of OntoMat-Service-Browser. Doing so, he could extend the web page with metadata if he has write access, primarily however he establishes mappings between concepts, relations and attributes from the ontology used by the web service provider to his client ontology [6].

In the following we describe the deep-annotation of the vendor web service shown in Figure 9. The web page advertising the web service describes the `getLaptopOffer` operation and constitutes the context for the usage of the vendor ontology. The aim of the annotator is to translate the terminology used in the description of `getLaptopOffer` (cf. the WSDL document in Figure 2 and the vendor ontology in Figure 4) into his client ontology (Figure 8).

```
:Product a rdfs:Class.
:id rdfs:domain :Product; rdfs:range :Literal.

:HardDisk a :Product.
:diskSize rdfs:domain :HardDisk; rdfs:range :Literal.
:Computer a :Product.
:hasHDD rdfs:domain :Computer; rdfs:range :HardDisk.
:price rdfs:domain :Computer; rdfs:range :Literal.
:cpuSpeed rdfs:domain :Computer; rdfs:range :Literal.

:Agent a :rdfs:Class.
:Company a :Agent.
:creditCardNumber rdfs:domain :Company; rdfs:range :Literal.

:Purchase a :rdfs:Class.
:hasBuyer rdfs:domain :Purchase; rdfs:range :Agent.
:hasObject rdfs:domain :Purchase; rdfs:range :Product.

:Insurance a :rdfs:Class.
:hasObject rdfs:domain :Insurance; rdfs:range :Product.
:price rdfs:domain :Insurance; rdfs:range :Literal.
:timePeriod rdfs:domain :Insurance; rdfs:range :Literal.
```

Figure 8: Ontology of the client

By drag'n'drop, one generates a graph of instances, relations between instances and attribute values of instances in the browser that visualizes the client ontology (cf. the left pane depicted in Figure 9).

When performing a drag'n'drop one will create a *literal instance*, if one drops

1. an instance of the vendor ontology onto a concept in the client ontology, or
2. a literal value onto a concept of the client ontology, or
3. if one drop's an attribute value of an instance onto an attribute in the client ontology.

For instance, dropping 'IBM' onto the concept company would create a corresponding literal instance in the client ontology, dropping '7MB' onto a size attribute of a selected instance creates a corresponding attribute value for this selected instance in the client ontology.

When performing a drag'n'drop one will create a *generic instance*, if one drops

- a concept A from the vendor ontology onto a client ontology concept B.

A generic instance is just a variable that states that concept A in the vendor ontology corresponds to concept B in the client ontology.[4]

Thus, one may augment the client ontology (represented in RDF by a graph G) by a graph G_Δ of new and different types of instances.[5] Each subgraph of G_Δ of non-separable, newly created instances and values in the client ontology corresponds to a mapping rule. For instance, one may *(i)*, drag'n'drop 'processorSpeed (from vendor ontology) onto cpuSpeed (from client ontology) that belongs to Computer (again in the client ontology). Thereby, *(ii)*, a generic instance is created for Computer with value Laptop (as cpuSpeed belongs to Computer and processorSpeed belongs to Laptop).

The corresponding interpretation in first-order logic is:

```
FORALL X (instanceOf(X,client:Computer) AND client:cpuSpeed(X,Y)) <-
        (instanceOf(X,vendor:Laptop) AND vendor:processorSpeed(X,Y)).
```

One may trace the later drag'n'drop action in Figure 9, where action 1 picks up 'Processor Speed' with its underlying web service parameter processorSpeed (cf. the markup `elementURI="http://laptop.wsdl/laptop/#part(getLaptopOfferRequest/processorSpeed)"` in Figure 7). It is dropped onto the attribute that comes closest in his client ontology, viz. the aforementioned cpuSpeed, and generates the consequences just mentioned. Similarly, the second text item "Disk Space" being annotated with the input parameter diskSpace is handled in action 2. This time, however, the annotator must also create a hasHDD relationship between the generic instance hardisk1 and the generic instance of computer1 to build a larger graph representing a mapping rule with two generic attribute values (on cpuSpeed and diskSpace). Finally, the annotator maps the output parameters in action 3 (cf. Figure 9).

[4]Corresponding generalizations exist for attributes and relationships.

[5]The newly populated ontology would then be $G' := G \cup G_\Delta$.

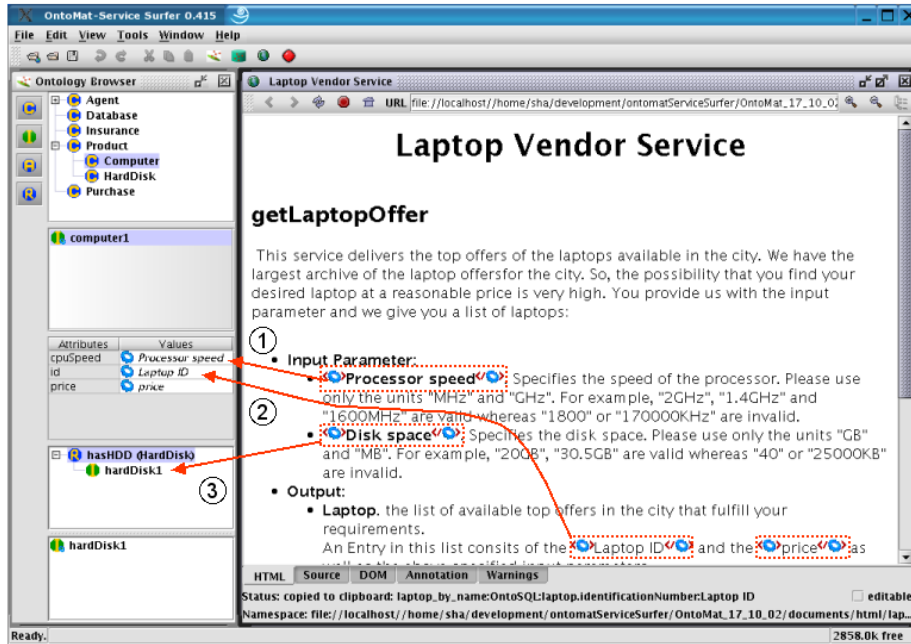

Figure 9: Screenshot of OntoMat-Service-Browser annotating vendor service

5.3 Investigating and Modifying Mapping Rules

The results of deep annotation are mapping rules between the client ontology and each service ontology. The annotator may publish the client ontology and the mapping rules derived from annotations. This enables third parties (in particular logical roles that follow in the OntoMat-Service process) to execute the services on the basis of the semantics defined in the client ontology.

We use F-Logic to define the mapping rules. F-logic is a deductive, object-oriented database language that combines the declarative semantics and expressiveness of deductive database languages with the rich data modelling capabilities supported by object oriented model [7].[6] However, the annotator does not have to write F-logic rules. They are generated automatically by the OntoMat-Service-Browser.

Figure 10 and Figure 11 give the reader an intuition of how such automatically generated mapping rules look like when visualized with the OntoEdit plugins OntoMap (cf., [6]). Figure 10 shows the mapping from the company ontology to the vendor ontology which is a result from the annotation effort indicated in Figure 9. The result for the corresponding mapping of the insurer's ontology is depicted in Figure 11.

[6]Thus in our implementation the aforementioned examplary mapping rule looks slightly different than the depicted first-order logic formulation. Since the first-order presentation is conceptually close enough, we have decided not to detract the reader by another syntax.

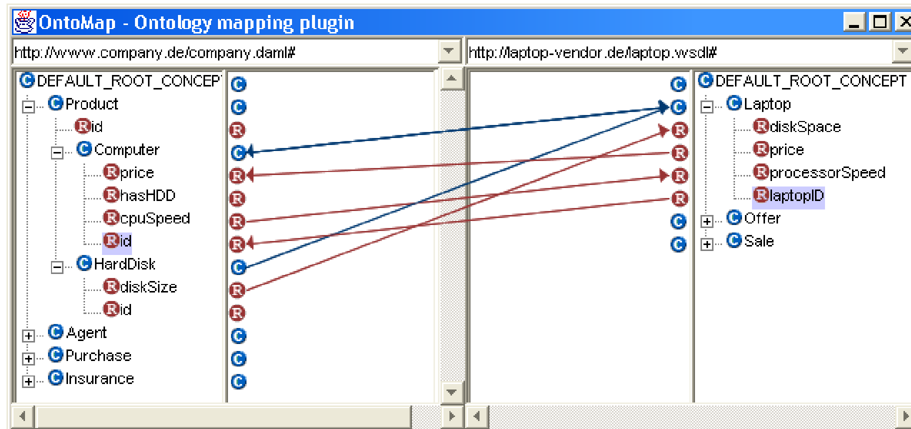

Figure 10: Mapping between Client Ontology (left window) and Vendor Ontology (right window)

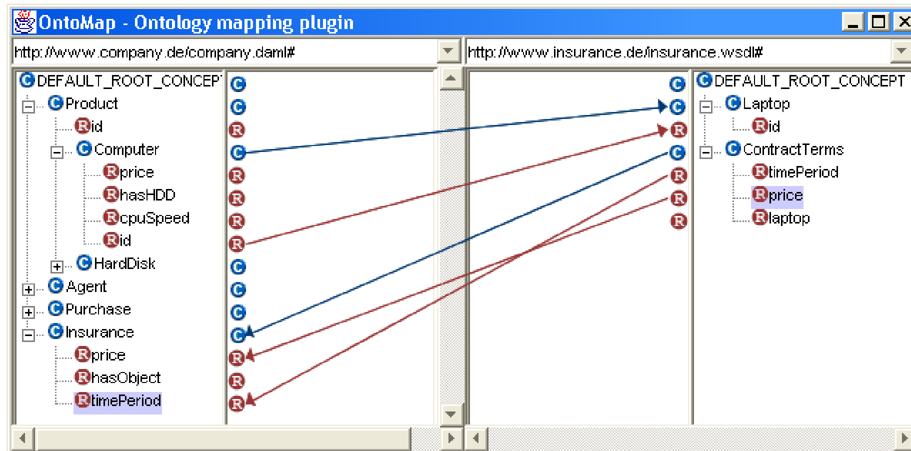

Figure 11: Mapping between Client Ontology (left window) and Insurer's Ontology (right window)

6 Web Services Planning

Often a user wishes to perform a task that is not directly accomplished by one single Web service. But in many cases there is a particular combination of Web services that would offer the needed functionality. Defining such a combination or composition of Web services manually from scratch can be difficult and time consuming especially for user with little technical background. Common AI planning techniques have been used in the past for performing such compositions automatically. However, the success of planning techniques is rather a disputed topic. We believe that AI planning techniques had moderate success in areas that tried to cover many different domains and aspects [14]. However, they were very successful in a small well defined domain with well defined building blocks [8]. Therefore, we believe that in this step of the OntoMat-Service, planning techniques can support a user in two ways[7]:

- Often a user needs a simple combination of Web services. With the help of planning techniques such combinations can be generated automatically. Thus, a user does not have to define the desired combination from scratch.
- If a user needs a complex composition to accomplish some task at hand, planning techniques help him by generating a rough "first version" of a combination (a plan), which the user can modify manually.

In Section 6.1 we show two alternative ways of specifying plans, viz. data-flow driven and control-flow driven. Section 6.2 uses the former to compute possible plans and Section 6.3 uses both types of plan specifications in order to present the generated results to the user. Thus, Section 6 proceeds through the third main step of the OntoMat-Service process.

6.1 Plan Specification

We use two paradigms to specify plans, namely *data-flow driven* and *control-flow driven*. The approaches are equivalent in so far as a given specification based on one paradigm can be translated into a corresponding specification based on the other paradigm. Each paradigm has its strengths and weaknesses. Depending on the context the one or the other plan specification approach should be preferred.

6.1.1 Data-Flow-Driven Plan Specification

Given a set of Web services W each web service $w \in W$ has a set of input parameters $w.I$ and a set of output parameters $w.O$. We define a connector $c = (o, i)$ with $o \in u.O$ and $i \in v.I$ and $u, v \in W$, when the output o of the Web service u becomes the input i of the Web Service v. The set of all such connectors is denoted by C. The set of Web services W together with a set of connectors C builds a directed data-flow graph, in which Web services in W build the vertices and the connectors in C build the edges.

Given such a data-flow graph, a Web service is executed as soon as the values of all its input parameters are available. This is the basic rule that determines the actual order of execution of the Web services.

[7] In [20] authors describe a composite Web service as a fixed template which must be configured for each specific use instead of pre- and postconditions based planning-style approach for composing a composite Web service from scratch everytime.

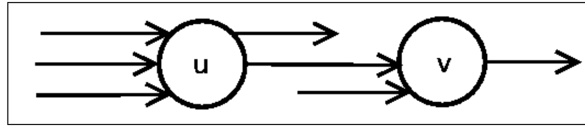

Figure 12: A data-flow graph that must be executed sequentially

Figure 12 shows a simple data-flow graph consisting of two Web services u and v. The arrows pointing in u and v represent the set of input parameters $u.I$ and $v.I$ respectively. The arrows pointing out of u and v represent the set of output parameters $u.O$ and $v.O$ respectively. Further, the arrow pointing out of u and into v connects the second input of u with the first input of v and represents a connector. Since the Web service v needs data from u, it must be executed after the Web service u.

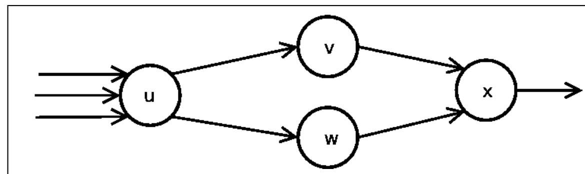

Figure 13: A data-flow graph that can be partially executed in parallel

Figure 13 shows another data-flow graph consisting of four Web services u, v, w and x. The arrows pointing in and out of a Web service and from a Web service to another Web service have the same meaning as in earlier example. Since the Web services v and w need data from u but not from each other, v and w must be executed after u but can be executed in parallel. Further, the Web service x must be executed after the completion of v and w since the values of its input parameters are available when both v and w are complete.

6.1.2 Control-Flow Driven Plan Specification

In contrast to the data-flow-driven approach where the order of execution is specified implicitly, in the control-flow driven approach the order of execution is specified explicitly with control constructs.

Currently, we support two types of control constructs, namely *sequence* and *parallel*. We denote the set of control constructs by S, that is, currently $C = \{sequence, parallel\}$[8]. With a component, we refer to either a Web service or a control-construct and denote the set of components by P, that is, $P = S \cup W$. A plan has exactly one main control-construct, which is executed when a plan is invoked. Now we describe the control-constructs and their execution semantics in more detail.

Sequence A sequence has an ordered set of components. We denote a sequence of components $p_1, \ldots p_n$ with $p_1, \ldots p_n \in P$ by $sequence(p_1, \ldots p_n)$. For a sequence s, we denote the set of its components with $s.P$. The execution

[8]In future, we will extend the set of supported control constructs by *choice, if-then-else, while, repeat-until* etc.

semantics of a sequence s is described recursively by the following rules

$$\begin{cases} \text{execute } p_1 \text{ then execute } \mathit{sequence}(p_2, \ldots, p_n) & \text{if } s.P \neq \emptyset \\ \text{do nothing} & \text{if } s.P = \emptyset \end{cases}$$

The plan in figure 12 can be specified in the control-flow-driven approach as $\mathit{sequence}(u, v)$.

Parallel Now we describe the control construct *parallel*. Like the sequence, the construct for parallelism p also has a set of components which we denote by $p.P$. Note, that $p.P$ does not need to be an ordered set. If the set of components $p_1, \ldots p_n \in P$ must be executed in parallel, then we denote it by $\mathit{parallel}(p_1, \ldots p_n)$. The execution semantics of a parallel construct p is described by the following rules

$$\begin{cases} \text{execute } p_1, p_2, \ldots, p_n \text{ in parallel} & \text{if } p.P \neq \emptyset \\ \text{do nothing} & \text{if } p.P = \emptyset \end{cases}$$

The execution of a construct for parallelism p is finished when the execution of all the elements of $p.P$ is finished. In case the actor executing the plan is a sequential actor, the construct for parallelism p can be implemented through a sequence with any permutation of the elements of $p.P$.

The plan in figure 13 can be specified in the control-flow driven approach as $\mathit{sequence}(u, \mathit{parallel}(v, w), x)$.

6.1.3 Integrating Mapping Rules in a Plan

If the output of a web service operation A is of type t and the input of another web service operation B is also of type t, then the service operations A and B can be plugged together (first A then B). Since, it is realistic to assume that different web service providers have different ontologies, this approach only support plans in which all the web services are provided by one web service provider or all the Web services providers refer to the same domain ontology. In the former case our mapping rules come into play. By using the mapping rules that align ontologies of different Web service providers, it is possible to deal with plans that contain Web services from different Web service providers. For example, if the output of a service A is of type t_1 and the input of another web service B is of type t_2 and there is a mapping rule from t_1 to t_2, the services A and B can be plugged together (first A then B).

Mappings are integrated in a plan by modelling them as a special kind of web services that are provided by the client himself. Currently, the premise as well as the conclusion part of our mapping rules is a conjunction of "instanceOf" terms. We interpret such a mapping by interpreting the terms in the premise of the rule as input parameters and the terms in the conclusion of the rule as output parameters of a Web service.

Given a set of such rules, OntoMat-Service-Browser automatically generates a set of corresponding Web services by interpreting the rules as described above. Consequently, the mapping rules are available to the user as Web services. These Web services can then be used in a combination of Web services just like other Web services.

6.1.4 Handling Multiple Occurences of Web Services

Above definition of a plan does not allow a web service to occur more than once in a plan. To deal with this we introduce the notion of a *web service occurence*. A web service occurence has a name and a reference to the web service it is an occurence of. The name should be chosen in a way such that it is unique among all the occurences of the web service in a plan, for example a running index. We denote an occurence of a web service w with name i by w^i. With $w^i.I$ we denote the set of input parameters and with w^i we denote the set of output parameter of the occurence i of a web service w. A plan would then contain web service occurences instead of web services.

6.2 Plan Generation

The planning component generates simple plans based on a given set of web services and a given set of mapping rules. The generated plans are specified by their data-flow graphs as described in Section 6.1.1.

The inputs and outputs of web services are specified in the web service description documents of the web services. By considering the mapping rules and the information about the input and output types of web services, the planning component is able to infer valid web service flows as follows.

The Web service end consumer selects the Web service, he wants to use to accomplish certain tasks at hand. By making such a selection, he restricts the sets of relevant mapping rules. The plan generation algorithm iterates over all selected Web services including the Web services that are interpretations of the relevant mapping rules and generates a data dependency graph (a directed acyclic graph). Figure 14 shows the data dependency graph that is generated if the user selects all four Web services that are mentioned in section 2 and assuming that there is a mapping rule between the concept `Laptop` in the vendor ontology and the concept `Laptop` in the insurance company ontology and this mapping rule is interpreted as Web service `m1`.

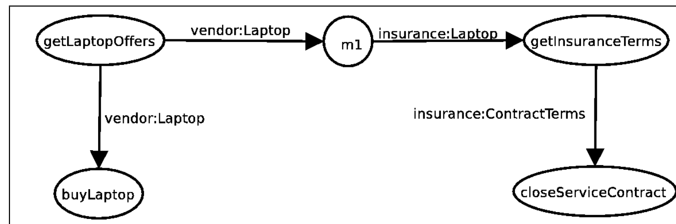

Figure 14: An example data dependency graph

As soon as the dependency graph is generated (polynomial time complexity in number of Web services), the user can define his goal by selecting the Web services whose outputs he is interested in. Starting with the goal Web services a set of subgraphs of the dependency graph is calculated by traversing the dependency graph backwards. Each such subgraph is the data-flow-driven specification of a plan whose execution would lead the user to his goal (cf. section 6.1).

In our running example, if the user selects only the Web service `Buylaptop` then the subgraph would contain the Web services `getLaptopOffers` and `buyLaptop`.

If the user selects the Web service `closeServiceContract` then the subgraph would contain the Web services `getLaptopOffers`, `m1`, `getInsuranceTerms` and `closeServiceContract`. If the user selects the Web services `buyLaptop` and `closeServiceContract`, then the subgraph is equal to the graph shown in figure 14.

6.3 Plan Presentation

Eventually, compiled web service plans are presented as part of the OntoMat-Service process. As elaborated on before, the inputs and outputs of a plan depend obviously on the inputs and outputs of the individual Web services, which are atomic from the client's point of view. For an individual Web service there is a Web page which the user can read to understand what the Web service does. But, there is no such Web page that describes a plan generated on the fly. Generating a descriptive Web page from the individual Web pages of the individual Web services is rather difficult, if not impossible.

To remedy the problem, we visualize the data-flow specification of a plan (cf. Figure 15). Taking advantage of the duality of data and control flow, we are currently implementing a presentation of the corresponding control-flow specification, too. The motivation is that the control-flow specification is frequently sparser than its data-flow counterpart — allowing a less crowded view on the overall plan and, thus, a better abstraction from too many details.

In addition, we present input and output parameters of (parts of) plans selected for investigation by the user. Inputs and outputs of a plan are calculated from the respective inputs and outputs of the individual Web services. Thereby, the set of input parameters of a plan is equal to the set of all the input parameters of all the individual Web services except those input parameters that are automatically available. Recall, that our set of connectors contain the information about the input parameters that are automatically available. Formally, the set of input parameters $p.I$ of a plan p is

$$p.I = \bigcup_{w \in W} w.I \setminus \{i, \text{ such that } \exists c = (o, i) \in C\}.$$

Similarly, the set of output parameters of a plan can be calculated as

$$p.O = \bigcup_{w \in W} w.O \setminus \{o, \text{ such that } \exists c = (o, i) \in C\}.$$

Note that $p.O$ does not contain the outputs that become inputs of subsequent web services. We believe that we need a more expressive specification of plans (e.g. one that can deal with messages and actors) to be able to handle the outputs that a user obtains during the execution of a plan.

7 Web Services Invocation

On the basis of the information about each plan, the user decides to execute a plan. Since OntoMat-Service-Browser currently does not generate arbitrary complex plans, we provide the user with the possibility to manually modify an automatically generated plan.

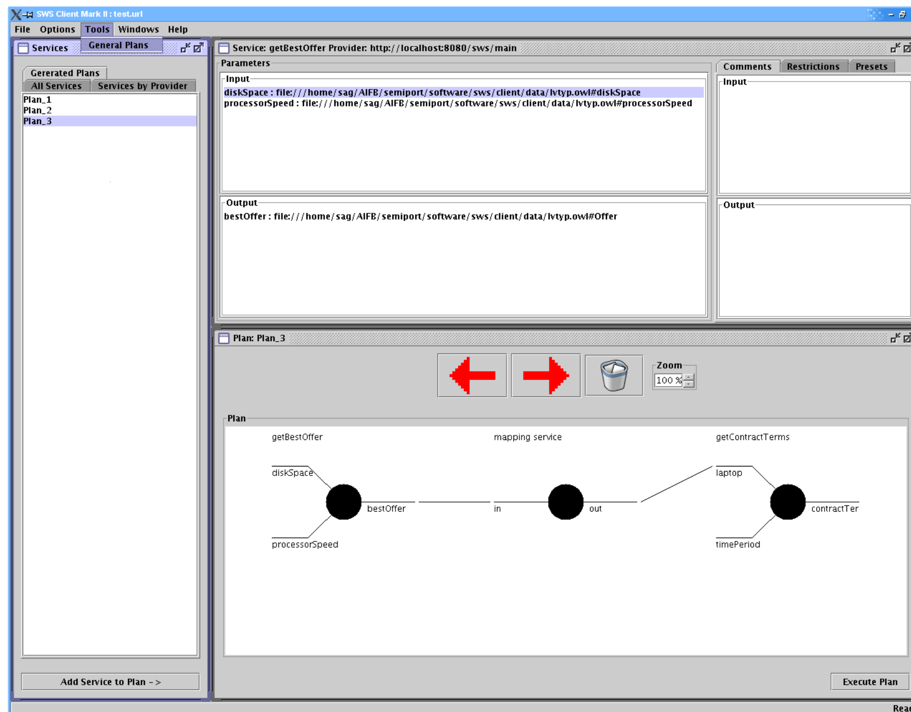

Figure 15: Plan Presentation as a Data-flow Graph

7.1 Configuring Access to Client's KB

Before web service execution begins, OntoMat-Service-Browser provides the user with a possibility to configure automatic retrieval of data that is needed during the execution. Depending on what a plan does, how long it is and how often it is executed by a user, this feature can be very helpful because it enables a certain degree of automation by preventing a user to re-enter the values for each input parameter manually during each instance of a plan. Obviously, this feature works only if the intended values of the input parameters are stored. For example, if a Web service asks for an email address of the user and the user has his email address stored in his local repository then he can configure that the value of the input parameter of the Web service should be retrieved from his local repository. Using this feature reduces the chances of unexpected behavior of a Web service since there are less typing mistakes. We specify one such configuration as the tuple `<uri,method,parameters>`, where `uri` represents the URI of the input parameter of the Web service the value of which should be automatically retrieved, `method` represents a programming language method that must be called in order to retrieve the data and `parameters` represent the set of parameters that must be passed to the `method`.

The actual invocation is performed by a generic web services client engine. Since we have implemented the invocation engine in Java, it can call methods of external Java classes. In this case, we describe the `method` part of the aforementioned tuple such that it points to a method of a Java class, which the invocation engine has access to.

7.2 Plan Execution and Generic User Interface

When the user requests the invocation of such a flow, the engine takes the plan, the set of mapping rules and the set of above mentioned configurations and calls the web services in the proper order. The order of execution of the Web services is implicitly given in the data dependency graph. A Web service is ready to be executed when the values of all its input parameters are available. The execution component communicates with OntoBroker [3], whenever mapping between concepts is required (cf. Figure 16) and calls the specified method whenever there is automatic retrieval configuration present for a required input parameter.

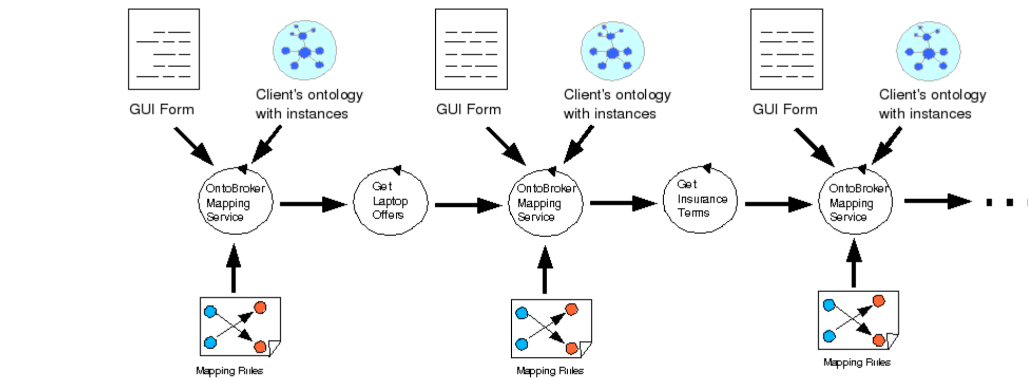

Figure 16: Service Flow in our Running Example

The invocation component differentiates between the following cases (cf. Figure 16):

- **There are no mapping rules:** In this case, the user is provided with a form like interface, in which he has to enter required data according to the ontology of the respective web service provider to proceed the execution (cf. figure 17).
- **Automatic retrieval of data from client's ontology is not configured and mapping rules are defined:** In this case, the user is provided with a form like user interface, in which he has to enter required data according to his own (client's) ontology.
- **Automatic retrieval of data from client's ontology is configured and mapping rules are defined:** In this case, the invocation runs fully automatically.

This kind of approach is a generalization of common approaches to invocation of single web service operations. Let us consider this simple case in our framework: If a user wants to manually call only one web service operation, he will skip the definition of mapping rules. The flow will consist of only one web service operation. When executing the single web service operation, the invocation engine will request data from the user via a form interface that reflects the ontology of the service provider (because no mapping rule exists).

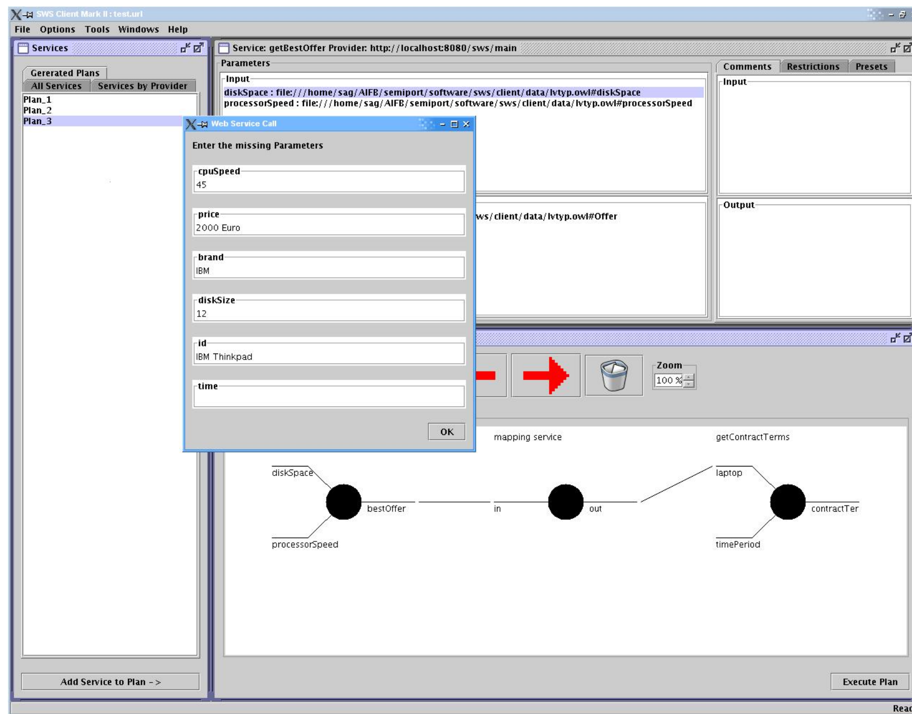

Figure 17: Example of generic user interface

8 Related Work

In this paper we provide an original framework, OntoMat-Service, to embed the process of web service discovery (here: by browsing web pages and retrieving web pages from search engines like Google), composition (here: by deep annotation and reasoning over logically possible configurations), and invocation (here: by OntoMat-Service-Browser, and the mapping to a client ontology). The consideration of semantic heterogeneity is germane to OntoMat-Service. It offers semantic translations as one of its core modules.

OntoMat-Service does not aim at a web service discovery, composition and invocation that is intelligent in the sense that it completely automates the task that typically the user is supposed to do. Rather, it provides an interface, OntoMat-Service-Browser, that supports the intelligence of the user and guides him to add semantic information such that only few logically valid paths remain to be chosen from by the user.

To fully pursue such an objective, one needs a large set of different modules. We have built on our existing experience and tool framework for semantic annotation (cf. [5, 6]) and for logical reasoning [3]. We have not yet dealt with the issue of web service flow execution and monitoring that is certainly needed to complement our current version of OntoMat-Service.

Closest to our approach come frameworks that facilitate the building of web service flows. A number of software systems are available to facilitate manual composition of programs, and more recently web services. Such programs,

which include a diversity of workflow tools [21, 4], and more recently service composition aids such as BizTalk Orchestration [10] enable a software engineer to manually specify a composition of programs to perform some task — though they typically neglect the aspect of semantic heterogeneity that is core to OntoMat-Service.[9]

Web Services Invocation Framework (WSIF) [18] is an open source framework to execute any web service, that can be described by a WSDL document. However, it does not support the execution of a flow of web services.

Some technologies have been proposed that use some form of semantic markup of web services in order to automatically compose web services to perform some desired task (e.g., [13, 2, 12]). In [13], the authors use situation calculus for representing web service description and Petri nets for describing the execution behaviors of web services. In [2], the authors present an architecture of intelligent brokers that offer problem solving methods that can be configured and used by the users according to their needs. In [12] the authors propose an extended version of Golog for formalizing the provision of high-level generic procedures and customization of constraints. In [17], the authors propose a rule based expert system to automatically compose web services from existing web services.

On the one hand most recent experiences from such advanced projects like IBrow, however, have shown that automatic composition techniques cannot yet been carried over to an open world setting. There one needs to tightly integrate the user of a web service — such as we do in OntoMat-Service. On the other hand OntoMat-Service can obviously be extended in the future to consider more types of automatic semantic matchmaking, service discovery [15, 19] and configuration of web services into the web service planning phase.

9 Discussion

In this paper we have described OntoMat-Service, an original framework to tie together the World Wide Web and web services into a Service Web. Germane to OntoMat-Service is its blending of browsing the Web, aggregating conceptual descriptions and web services and then investigating and invoking them from one platform.

We have also presented OntoMat-Service-Browser, a tool that constitutes a prototype implementation of OntoMat-Service. Currently, our prototype understands WSDL with RDF(S) for web service descriptions, but its flexible architecture allows easy integration of more powerful web service description languages like DAML-S [1].

Clearly, one must be aware of what OntoMat-Service and OntoMat-Service-Browser can do and what they can't do. OntoMat-Service is not intended to cater to businesses that want to establish complex web service connections with intricate interactions. For this objective, the integration by semantic annotation may provide a quick, first prototype, but semantic annotation cannot provide arbitrary complex mapping rules or arbitrarily complex workflows. On the other hand, OntoMat-Service allows exactly for easily building a prototype web service integration and it allows for users with domain knowledge (e.g. consultants

[9]BizTalk even allows for XML-based (non-semantic) translations of data.

doing ERP configuration) to participate in the Service Web — without much programming.

OntoMat-Service opens up many interesting questions that need to be solved in the future, such as

- how to automate the way that Web Services are presented to the World;
- how to characterize the boundaries of what functionality can be aggregated and executed.
- how to annotate mappings between ontologies (semi-) automatically [16].

Eventually, OntoMat-Service and OntoMat-Service-Browser, in conjunction with their counterparts in semantic annotation [5] and deep annotation [5], open up the possibility to bring Web pages, databases and Web Services into one coherent framework and thus progress the Semantic Web to a large Web of data and services.

10 Acknowledgements

Part of this work was funded by the BMBF (federal ministry of education and research) projects SemIPort and internetökonomie (SESAM).

References

[1] Anupriya Ankolekar, Mark Burstein, Jerry R. Hobbs, Ora Lassila, Drew McDermott, David Martin, Sheila A. McIlraith, Srini Narayanan, Massimo Paolucci, and et al. Terry Payne. Daml-s: Web service description for the semantic web. In *1st Int'l Semantic Web Conf. (ISWC 02)*, 2002.

[2] V. Richard Benjamins, Enric Plaza, Enrico Motta, Dieter Fensel, Rudi Studer, Bob Wielinga, Guus Schreiber, and Zdenek Zdrahal. Ibrow3 - an intelligent brokering service for knowledge-component reuse on the world wide web. In *Proc.11th Banff Knowledge Acquisition for Knowledge-Based System Workshop (KAW98)*, 1998.

[3] S. Decker, M. Erdmann, D. Fensel, and R. Studer. Ontobroker: Ontology based access to distributed and semi-structured information. In *DS-8 — Proceedings of the Conference on Database Semantics*, pages 351–369, 1999.

[4] C.A. Ellis and G.J. Nutt. Modelling and enactment of workflow systems. *Application and Theory of Petri Nets*, LNCS 691:Modelling and enactment of workflow systems, 1993.

[5] S. Handschuh and S. Staab. Authoring and annotation of web pages in cream. In *Proceedings of the 11th International World Wide Web Conference, WWW 2002, Honolulu, Hawaii, May 7-11, 2002*, pages 462–473. ACM Press, 2002.

[6] S. Handschuh, S. Staab, and R. Volz. On deep annotation. In *Proceedings of the 12th International World Wide Web Conference, WWW 2003, Budapest, Hungary, May 20-24, 2003 (to appear)*. ACM Press, 2003.

[7] Michael Kiefer, Georg Lausen, and James wu. Logical foundations of object oriented and frame-based languages. *Journal of the ACM*, 1995.

[8] Jana Koehler and Kilian Schuster. Elevator control as a planning problem. In *Artificial Intelligence Planning Systems*, pages 331–338, 2000.

[9] O. Lassila and R. Swick. Resource description framework (RDF) model and syntax specification. Technical report, W3C, 1999. W3C Recommendation. http://www.w3.org/TR/REC-rdf-syntax.

[10] David Lowe, Xin Chen, Todd Mondor, Tomislav Rus, Ned Rynearson, Steve Wright, and Tom Xu. *BizTalk(TM) Server: The Complete Reference.*, November 2001.

[11] Alexander Maedche and Steffen Staab. Services on the move — Towards p2p-enabled semantic web services. In *Proceedings of the 10th International Conference on Information Technology and Travel & Tourism, ENTER 2003, Helsinki, Finland, 29th-31st January 2003*. Springer, 2003.

[12] S. McIlraith and T. Son. Adapting golog for composition of semantic web services. In *Proc 8th International Conference on Principles of Knowledge Representation and Reasoning*, 2002.

[13] Srini Narayanan and Sheila McIlraith. Simulation, verification and automated composition of web services. In *WWW2002*, May 2002.

[14] Nils J. Nilsson. Shakey the robot. Technical Report 323, AI Center, SRI International, 333 Ravenswood Ave., Menlo Park, CA 94025, Apr 1984. http://www.sri.com/about/timeline/shakey.html.

[15] M. Paolucci, T. Kawmura, T. Payne, and K. Sycara. Semantic matching of web services capabilities. In *First Int. Semantic Web Conf.*, 2002.

[16] Abhijit Patil, Swapna Oundhakar, Amit Sheth, and Kunal Verma. Meteor-s web services annotation framework. In *Proc. of the Thirteenth International World Wide Web Conference (WWW 2004)*, 2004.

[17] Shankar R. Ponnekanti and Armando Fox. Sword: A developer toolkit for building composite web services. In *Proceedings of the 11th International World Wide Web Conference, WWW 2002, Honolulu, Hawaii, May 7-11, 2002*. ACM Press, 2002.

[18] Apache Web Services Project. Web services invocation framework. http://ws.apache.org/wsif.

[19] Katia P. Sycara, Matthias Klusch, Seth Widoff, and Jianguo Lu. Dynamic service matchmaking among agents in open information environments. *SIGMOD Record*, 28(1):47–53, 1999.

[20] Annette ten Teije, Frank van Harmelen, and Bob Wielinga. Configuration of web services as parametric design. In *Proceedings of the Proceedings of the 14th International Conference on Knowledge Engineering and Knowledge Management (EKAW'04)*, Lecture Notes in Artificial Intelligence. Springer-Verlag, 2004.

[21] van der Aalst W.M.P. Woflan: A petri-net-based workflow analyzer, systems analysis - modelling - simulation. *Systems Analysis - Modelling and Simulation*, 35(3):345–357, 1999.

[22] W3C. Web service description language (wsdl) version 1.2, March 2003. http://www.w3.org/TR/wsdl.