

The PORSCE II Framework: Using AI Planning for Automated Semantic Web Service Composition

OURANIA HATZI¹, DIMITRIS VRAKAS², NICK BASSILIADES², DIMOSTHENIS ANAGNOSTOPOULOS¹ and IOANNIS VLAHAVAS²

¹*Department of Informatics and Telematics, Harokopio University of Athens, Harokopou 89, 17671, Athens, GREECE*
E-mail: {raniah,dimosthe}@hua.gr

²*Department of Informatics, Aristotle University of Thessaloniki, 54124, Thessaloniki, GREECE*
E-mail: {dvrakas,nbassili,vlahavas}@csd.auth.gr

Abstract

This paper presents PORSCE II, an integrated system that performs automatic semantic web service composition exploiting AI techniques, specifically planning. Essential steps in achieving web service composition include the translation of the web service composition problem into a solver-ready planning domain and problem, followed by the acquisition of solutions, and the translation of the solutions back to web service terms. The solutions to the problem, that is, the descriptions of the desired composite service, are obtained by means of external domain-independent planning systems, they are visualized and finally evaluated. Throughout the entire process, the system exploits semantic information extracted from the semantic descriptions of the available web services and the corresponding ontologies, in order to perform composition under semantic awareness and relaxation.

1 Introduction

The World Wide Web has evolved from a collection of documents into a more integrated environment, where not only information but also system functionality is exposed and the provision of services plays an important role. The web service technology is a fundamental part of the web, as it provides a standard way to interact with information systems, independent from platform and internal implementation, thus accommodating interoperability between heterogeneous systems. However, in many cases, the need for integrated functionality cannot be fulfilled by a simple atomic web service, leading to the requirement for web service composition; that is, the appropriate combination of atomic web services in order to achieve a complex goal. The task of web service composition becomes significantly difficult, time-consuming and inefficient as the number of available atomic services increases continuously; therefore, the possibility to automate the web service composition process is proved essential.

Automated web service composition is significantly facilitated by the development of the Semantic Web, which permits the representation of knowledge about the actual meaning of information and services. The existence of such semantic information enables composition using intelligent techniques, such as AI Planning. Without the presence of semantic information, a high degree of human expertise would be required in order to compose web services meaningfully and not based on circumstantial syntactic similarities. The incorporation of semantics in the description of web services is accommodated through the development of a number of standard languages such as OWL-S (OWL-S) and SAWSDL (SAWSDL). Nevertheless, there are no tools utilizing semantic information incorporated in OWL-S to efficiently compose web services either accurately or approximately, taking into account the actual meaning of web service inputs/outputs as well as the corresponding ontologies.

The PORSCE II framework aims at automated semantic web service composition by employing planning under semantic awareness and relaxation. Its contribution focuses on the effective utilization

of semantic information present in OWL-S description of web services to enhance the web service composition process by facilitating approximate compositions, via planning.

The first and decisive step in the process concerns the translation of the web service composition problem to AI planning terms (Hatzi et al., 2009). This translation is based on the observation that certain correspondences exist between the two domains, which, given the appropriate mapping, enable the application of planning techniques to solve the web service composition problem effectively. Such correspondences include the available web services that can be combined to formulate meaningful compositions, which can be mapped to the planning domain, and user requirements about the desired composite service, which can be perceived as a planning problem over this domain.

The translation takes place between the most prominent standards in each area: OWL-S for semantic description of web services (either atomic or composite) and PDDL (Planning Domain Definition Language) (Ghallab et al., 1998) for definition of planning domain and problem. According to user preferences, the translation process may take into account semantics, resulting from the semantic analysis of the domain and the corresponding ontologies; if so, semantically equivalent or relevant concepts are also included, in order to cope with cases when no exact plans can be found and approximations must take place. The result of this phase of the transformation process is a fully formulated, solver-ready planning problem which incorporates all the required semantic information. PORSCE II consequently exports the planning problem to PDDL and invokes external planning systems to acquire plans, which constitute descriptions of the desired composite service. Each composite service is evaluated in terms of statistic and accuracy measures, while a visual component is also integrated, which accommodates composite service visualization and manipulation. Modification in the composite service is performed by atomic service replacement, either with an alternative equivalent atomic service, or through finding a sub-plan that can substitute it. If necessary, the composite service can also be modified through replanning. Finally, in order to provide full-cycle support, and render the result of the composition process independent from planning, the composite service is translated back to OWL-S, presenting the user with a description in the same standard as the initial atomic services and facilitating composite service deployment.

The rest of the paper is organized as follows: Section 2 discusses related work in the area of web service composition through planning, while Section 3 provides background related to the OWL-S standard, AI planning and the PDDL standard. Section 4 outlines system architecture, Section 5 elaborates on the main knowledge engineering aspect that this paper focuses on, that is, the transformation process, and Section 6 presents the rest of the system operations. Section 7 presents a case study and performance evaluation results and finally, Section 8 concludes and poses future directions.

2 Related Work

A number of approaches for automatic web service composition can be found at (Rao and Su, 2004) and (Dustdar and Schreiner, 2005); the most closely related with the approach proposed in this paper are discussed and summarized in Table 1.

SHOP-2 (Sirin et al., 2004) uses services descriptions in DAML-S, the predecessor of OWL-S, and performs Hierarchical Task Network (HTN) planning to solve the problem. The main disadvantage of this approach lies in the fact that the planning process, due to its hierarchical nature, requires the specification of certain decomposition rules, which have to be encoded in advance by an expert in the specific domain, with the help of a DAML-S process ontology.

OWLS-XPlan (Klusch and Gerber, 2005) uses the semantic descriptions of atomic web services in OWL-S to derive planning domains and problems, and invokes a planning module called XPlan to generate the composite services. The system is PDDL compliant, as the authors have developed an XML dialect of PDDL called PDDXML. Although the system imports semantic descriptions, the semantic information provided from domain ontologies is not utilized and semantic awareness is not achieved; therefore the planning module requires exact matching for service inputs and outputs.

The SWORD framework (Ponnekanti and Fox, 2002) is a set of tools that addresses a particular subset of the automatic web service composition problem. SWORD follows an approach close to the one of PORSCE II for modelling web service composition. However, atomic services are represented as rules,

stating that given specific inputs, an atomic service produces specific outputs. SWORD then uses a rule-based expert system to generate the composition plan, rather than classical planning.

The work described in (McIlraith and Son, 2002) conceives web service composition as a planning and execution task, where actions (services) may be complex. The GOLOG language is then employed, adapted and extended to address the issue of composition. Approaches that use knowledge-based planning include the work in (Pistore et al., 2005). Web service descriptions are expressed in a standard process modeling and execution language, such as BPEL4WS, therefore some prior, domain-specific knowledge of the composition issues is required, while another approach employing estimated-regression planning is presented in (Mcdermott, 2002); however, in order to be used, it requires extension to current standards.

System / Approach	Atomic Services	Composite Service (Inputs / Outputs)	Advantages	Disadvantages
SHOP-2	Primitive Tasks	Task Networks, Compound Tasks	Heuristics for increased performance	Prior expert domain knowledge required
OWLS-XPlan	Primitive Tasks	Initial State / Goal State	Combination of planning techniques	Ontology information not utilized. Domain-specific knowledge required
SWORD	Entity-Relation Model, Horn r.	Initial State / Goal State	Utilizes research in the area of rule-based expert systems	Requires user intervention. Not straightforward representation. Requires domain-specific knowledge
GOLOG	Situation Calculus	High-level generic procedures and constraints	Utilizes research in the area of situation calculus	Complex encoding and translation. Decreased scalability and interoperability
Estimated-Regression Planning	State Trans. Ops	Initial State / Goal State	Heuristics for acceleration of the process	Requires extension to the standards and planners. Decreased scalability
Knowledge-based	BPEL Processes	Compound process - goal	Well founded formalization. Non-determinism.	Semantic information and ontologies not utilized. Decreased scalability

Table 1 Overview and comparison of planning approaches to web service composition.

The main advantages of the proposed framework with respect to the aforementioned systems include the extended utilization of semantic information, in order to perform planning under semantic awareness and relaxation, and find better and, when necessary, approximate solutions. Furthermore, PORSCE II requires neither prior domain-specific knowledge, nor any kind of extension to the standards, in order to form valid, desired composite services; the OWL-S descriptions of the atomic web services and the corresponding ontologies suffice. Finally, PORSCE II is able to scale up for a great number of services, having the flexibility to exploit modern, advanced planners, and it also handles cases of service failure or unavailability dynamically, an important feature not covered by the aforementioned frameworks.

3 Background

This section presents the fundamental standards for the proposed approach, namely OWL-S and PDDL. Additionally, it introduces some basic planning notation that will be used throughout the paper.

3.1 OWL-S

OWL-S is an upper ontology based on OWL (OWL), created in the context of the Semantic Web in order to describe knowledge concerning semantic web services. It is used in combination with ontologies organizing the concepts appearing in the OWL-S descriptions. The use of OWL-S renders the semantics of the descriptions machine comprehensible; therefore it enables intelligent agents to discover, invoke and compose web services automatically. A web service description in OWL-S is comprised of (OWL-S):

- *Service Profile*: describes what the service accomplishes, limitations on service applicability and quality, and requirements that the service requester must satisfy to use the service.
- *Process Model*: describes the way a client can communicate and use the service.

- *Service Grounding*: specifies concrete details of how an agent can access a service, such as communication protocols and message formats.

The proposed approach utilizes semantic information contained in the *Service Profile*, along with the corresponding ontologies, in order to translate the description in planning terms.

An ontology in this context refers to a formal representation of the concepts appearing as inputs and outputs in the web service profiles. The concepts in the ontology are connected with hierarchical relationships, such as superclass, subclass and sibling.

Apart from atomic web services, which involve atomic processes, OWL-S establishes a framework for defining composite processes as well. A composite process consists of a set of atomic processes, combined together using a number of control constructs, such as Sequence, Split, Split+Join, Choice, Any-Order, Condition, If-Then-Else, Iterate, Repeat-While, and Repeat-Until. The main reasons for using these constructs while defining a composite web service are: a) to enable the definition of compact services (e.g. through the use of Iterate, Repeat-While and Repeat-Until), b) to facilitate the definition of alternative paths (i. e. through the use of Conditions and If-Then-Else constructs) and c) to speed up the invocation of the composite web service, by allowing multiple atomic processes to be invoked concurrently (i.e. through the use of Split and Split+Join constructs).

3.2 Planning & PDDL

A planning domain and problem is usually modeled according to STRIPS (Stanford Research Institute Planning System) notation (Fikes and Nilsson, 1971) as a tuple $\langle I, A, G \rangle$, where I is the initial state, A is a set of available actions and G is a set of goals. States in STRIPS are represented as sets of atomic facts. Set A contains all the actions that can be used to modify states. Each action A_i has three lists of facts containing the preconditions of A_i , the facts that are added, and the facts that are deleted from the world state after the application of the action, noted as $prec(A_i)$, $add(A_i)$ and $del(A_i)$ respectively.

The following formulae hold for the states in the STRIPS notation:

- An action A_i is applicable to a state S if $prec(A_i) \subseteq S$.
- If A_i is applied to S , the successor state S' is calculated as $S' = S - del(A_i) \cup add(A_i)$
- The solution to a planning problem (plan P) is a sequence of actions $P = A_1, A_2, \dots, A_n$, which, if applied to I , lead to a state S' such that $S' \supseteq G$.

Planning Domain Definition Language (PDDL) (Ghallab et al., 1998) was initially designed for providing a standard means of encoding planning domains and corresponding problems used as input test sets for planners that took part in planning competitions such as IPC (IPC2004). However, it has since been enhanced, extended and become a standard for modeling planning domains and problems.

PDDL provides structures to represent all the aforementioned STRIPS elements, such as predicates (atomic facts), actions and problems. Newer versions of the language (Gerevini and Long, 2005) added more features in order to enable the representation of more complex domains. These features include constants, variables, functions, and numeric expressions. PDDL also provides separate structures that can be used to represent problems, which are associated with specific planning domains. The latest extensions to the PDDL standard take into account the temporal properties of domains, as well as quality metrics, features which might prove very useful in the web service composition case, while PDDL+ (Fox and Long, 2002), also provides a standard way to represent plans, either sequential or partially parallel.

4 Overview and Architecture

PORSCE II was based on the results obtained from the prototype system PORSCE (Hatzi et al., 2008). PORSCE II aims at a high degree of integration as, along with the core transformation component, it additionally contains a visual interface, more elaborate relevance metrics, the ability for composite service accuracy assessment, and composite web service manipulation features. Furthermore, PORSCE II adopts a way of modeling the web service composition as a planning problem which reduces the complexity of the generated planning problem, thus accelerating the planning process. In order to highlight the planner independency of PORSCE II, which enables the use of any domain independent planning system based

on PDDL, two external planners have been included. Finally, PORSCE II supports seamless composition, by initiating the process with the OWL-S descriptions of atomic web services, and concluding with the OWL-S description of the produced composite service, thus facilitating deployment.

The key features of the framework are:

- Translation of OWL-S web service descriptions (atomic or composite) into planning operators.
- Interaction with the user in order to acquire their preferences regarding the desired composite service and desired metrics for semantic relaxation.
- Enhancing the planning domain and problem with semantically similar concepts.
- Exporting the web service composition problem as a PDDL planning domain and problem.
- Acquisition of solutions by invoking external planners.
- Flexibility in the choice of planner, as any PDDL-compliant external planning system can be used.
- Assessing the accuracy of the composite services.
- Visualizing and modifying the solution by atomic service replacement or replanning.
- Transformation of the solution (composite web service) back to OWL-S.

PORSCE II comprises of the OWL-S Parser, the Transformation Component, the OWL Ontology Manager, the Visualizer and the Service Replacement Component. An overview of the architecture and the interactions among the components is depicted in Fig. 1.

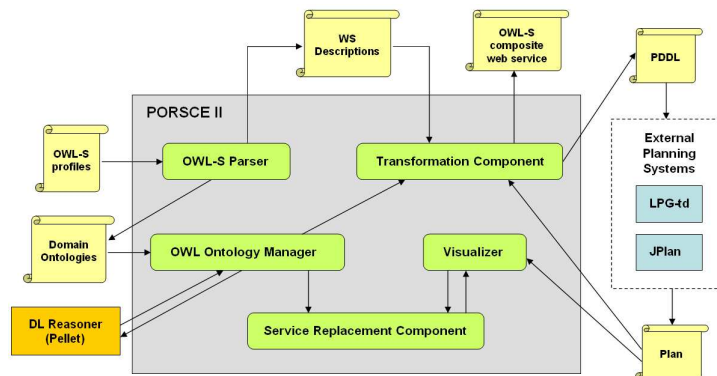


Figure 1 PORSCE II architecture.

The OWL-S Parser is responsible for parsing a set of OWL-S web service profiles and determining the corresponding ontologies that the concepts appearing in the web service descriptions belong to. The OWL Ontology Manager (OOM), utilizing the inferencing capabilities of the Pellet DL Reasoner (Sirin et al., 2007), applies the selected algorithm for discovering concepts that are similar to a query concept. The Transformation Component is responsible for a number of operations that result in the formulation of the planning problem from the initial web service composition problem, its consequent solving, and the transformation of the produced composite service back to OWL-S. The purpose of the Visualizer is to provide the user with a visual representation of the plan, which in fact is the description of the composite service. Finally, the Service Replacement Component enables the user to employ a number of alternative techniques in order to replace a specific atomic web service in the composite service sequence. PORSCE II is implemented in Java and it is available online, along with example problems, at <http://www.dit.hua.gr/~raniah/porsceII.en.html>.

5 Transformation Process

The transformation process includes translation of the web service composition problem into a planning problem and possible enhancement with semantic information, as well as transformation of the plan representing the produced composite service back in web service context. The process starts at the OWL-S Parser, which parses the OWL-S profiles of the available atomic web services and forwards them to the Transformation Component. The Transformation Component is responsible for a number of operations, including translating the web service descriptions received from the OWL-S Parser to planning operators and enhancing them with similar concepts derived from the OOM. Moreover, it interacts with the user in

order to formulate the planning problem, and exports both the planning domain and problem to PDDL. Finally, it translates the PDDL+ plan back to OWL-S, completing the composition process.

5.1 OWL-S to PDDL Translation

The first step in the translation process generates the planning domain by translating each available OWL-S web service profile WSD_i into a planning action A_i (Fig. 2). More specifically:

- The name of the action is the *rdf:ID* field of the profile: $name(A_i) = WSD_i.ID$
- The preconditions of the action are formed by the service input and precondition definitions:

$$prec(A_i) \equiv \bigcup_{k=1}^n WSD_i.hasInput_k \cup \bigcup_{k=1}^m WSD_i.hasPrecondition_k$$

- The add effects of the action comprise of the service output and positive effect definitions:

$$add(A_i) \equiv \bigcup_{k=1}^n WSD_i.hasOutput_k \cup \bigcup_{k=1}^m WSD_i.hasEffect_k^+$$

- The delete list is formed by the negative effect definitions. The SWRL language (SWRL) was used in order to model the preconditions and effects of the web services. SWRL combines OWL DL and RuleML in order to model preconditions and consequences in the Semantic Web, through the use of Horn-like rules. In the PORSCHE II case, preconditions are modeled by SWRL rule conditions, while positive effects are modeled as SWRL atomic expressions that are true in the world after the execution of the web service. Since SWRL does not directly support for negation and negated atomic expressions, which would model negative (delete) effects, the negation `<neg>` element of RuleML (RuleML) was employed, which is used by the transformation process in order to discriminate between add and delete effects. The delete list of the action is formed as follows:

$$del(A_i) \equiv \bigcup_{k=1}^m WSD_i.hasEffect_k^-$$

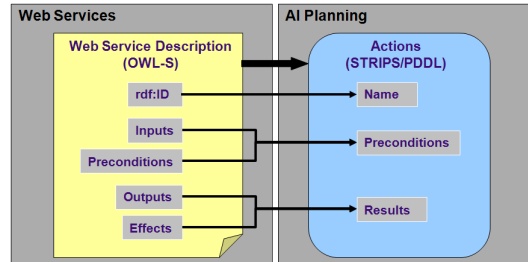


Figure 2 Web services to planning domain mapping.

This transformation can be applied either to atomic or to composite web services described in OWL-S, provided that the outputs of the composite service are fully deterministic; PORSCHE II is not concerned with the internal implementation of the services. An example of an OWL-S to PDDL transformation is presented in Fig. 3, where the mapping presented above is marked. The web service description at hand concerns a web service that accepts as inputs the ISBN, a Seller and a Client of a book, and has as a precondition that a Seller possesses the book indicated by the ISBN. The service has as output the Order Data, and as a result the change of book ownership from Seller to Client. The addition of the auxiliary argument in every PDDL predicate representing an input or output concept was a necessary technicality, in order to overcome the disadvantage of many planners that did not accept predicates without arguments, which does not in any way affect the outcome of the composition.

After the creation of the planning domain, the next step is the generation of a corresponding planning problem, based on the user requirements about the composite service. A straightforward solution adopted by PORSCHE II for this step is the following: Let IC be the set of concepts that the user wishes to provide to the composite service and GC its desired outputs. If O denotes the set of all available concepts in the ontology, then $IC \subseteq O$, $GC \subseteq O$ and $IC \cap GC \equiv \emptyset$. The inputs that the user wishes to provide formulate the initial state of the planning problem, while the desired outputs of the composite service formulate the goals: $I = IC$ and $G = GC$. Both input and output sets are provided externally by the user.

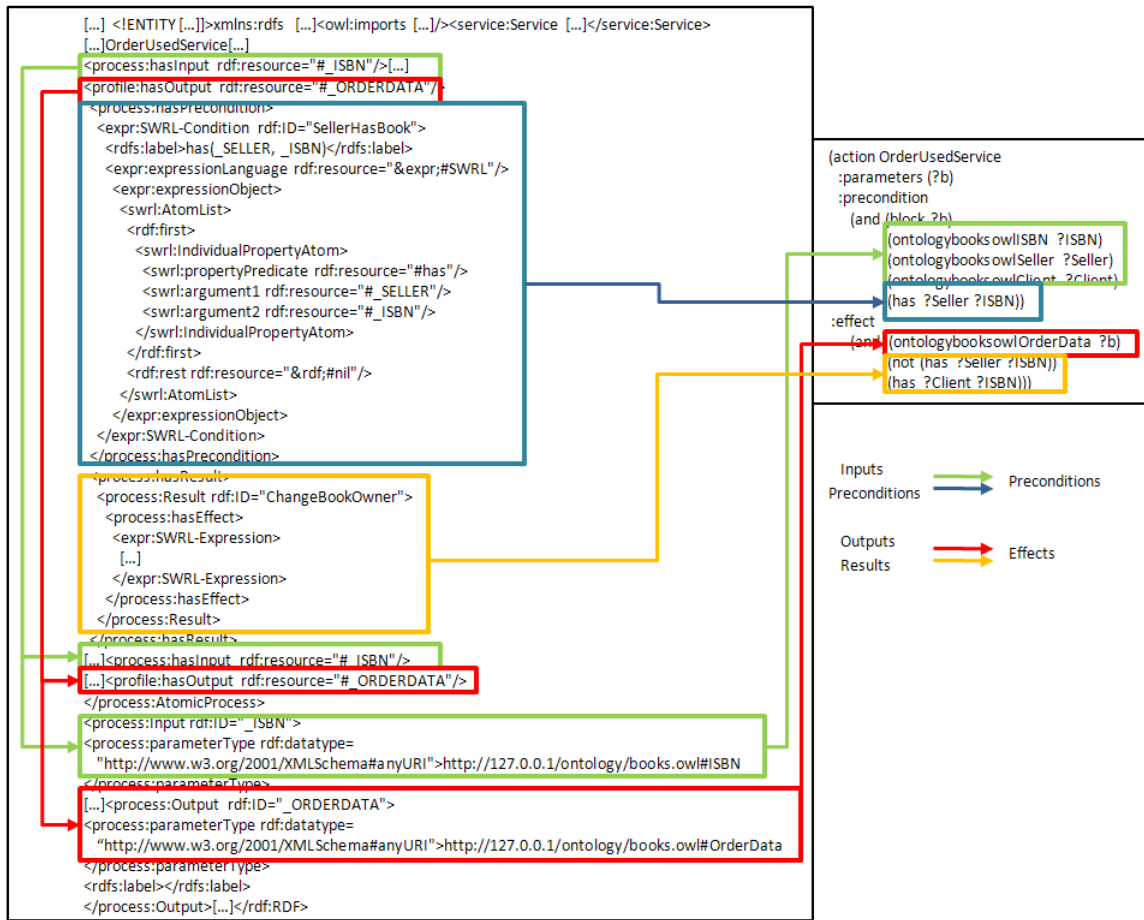


Figure 3 OWL-S to PDDL translation example.

5.2 Semantic Analysis

The step of semantic analysis, following that of the transformation process described in the previous subsection, enables the system to exploit semantic information. This step is implemented by the OWL Ontology Manager (OOM). During translation, the OOM is used extensively for performing semantic relaxation, which is useful in cases when an exact input-to-output matching plan is not available. The OOM locates equivalent and semantically relevant concepts; therefore, approximate plans can be created.

In our approach, two ontology concepts are considered semantically similar if and only if they have a *hierarchical relationship* and their *semantic distance* does not exceed a user-defined threshold.

As far as the *hierarchical relationship* is concerned, four hierarchical filters are used for its definition for two ontology concepts A and B:

- *exact*(A, B): The two concepts should have the same identifier (URI) or they should be equivalent, in terms of OWL class equivalence, i.e. $A = B \vee A \equiv B$.
- *plugin*(A, B): The concept A should be subsumed by the concept B, i.e. $A \sqsubseteq B$.
- *subsume*(A, B): The concept A should subsume the concept B, i.e. $B \sqsubseteq A$. In both the *plugin* and the *subsume* filters the subsumption relationships of equivalent concepts are not considered.
- *sibling*(A, B): The two concepts should neither have a hierarchical relationship, nor be disjoint; instead, they should have a common superclass T, such as $A \sqsubseteq T \wedge B \sqsubseteq T$.

The *semantic distance* between two ontology concepts is calculated in PORSCE II using two methods:

1. The *Edge-Counting Distance* (*ec*) is based on the observation that in the hierarchical, tree-like structure of an ontology, the further two concepts are placed, the less semantically related they are. Therefore, it computes the semantic distance of two concepts in terms of the number of edges (*p*) found on the shortest path between them in the ontology tree. An edge exists between two concepts

A and B if A is the direct subclass of B , denoted as $A \sqsubseteq_d B$. The implementation of the *ec* distance between two concepts, denoted as $d_{ec}(A, B)$, returns a value between 0 and 1, with 1 denoting absolute mismatch. This value is the result of normalization of the number of edges to $[0..1]$ as p/p_{max} , using the maximum *ec* distance (p_{max}) found in the ontology. For performance purposes, p_{max} can be approximated as $p_{max} = 2h - 1$, where h is the maximum number of edges between the root owl:Thing concept (T) and the furthest leaf.

2. The *Upwards Cotopic Distance*, denoted as $d_{uc}(A, B)$, is defined in terms of the upwards cotopic measure, denoted as $uc(A)$ that represents the set of the superclasses of the concept A , including A itself (Maedche and Zacharias, 2002). In PORSCCE II, the upwards cotopic distance definition has been modified to incorporate the semantics of an ontology hierarchy. More specifically, the owl:Thing concept is not considered in the *uc* measure, while the union and intersection set operators take into account the concept equivalence semantics, thus ignoring concept set multiplicity. The upwards cotopic distance is defined as

$$d_{uc}(A, B) = 1 - \frac{|uc(A) \cap uc(B)| - 1}{|uc(A) \cup uc(B)| - 1}$$

If two concepts are disjoint, then their distance equals 1; otherwise, if the two concepts have a hierarchical relationship, then $d_{uc}(A, B) \in [0..1]$. The upwards cotopic measure reflects the significance of the common ancestors of two concepts in the ontology hierarchy, based on the intuition that concepts with a great fraction of common ancestors among all their ancestors tend to be semantically related.

5.3 Semantic Awareness and Relaxation

After the steps of the translation process and the semantic analysis are both complete, the system is able to perform semantic awareness and semantic relaxation. This step is potentially essential, because the representation of the web service composition as a planning problem is significantly empowered if the planning system is aware of semantic similarities among syntactically different concepts.

The implementation in PORSCCE II involves enhancing the domain and problem description with all the required semantic information and consequently letting the planner handle it as a classical planning problem. This solution is employed in order to: a) be able to use any PDDL compliant planner, as the semantic enhancement applied to the domain remains transparent to the planner, and b) minimize the interactions between the planner and the OOM, which introduce an overhead on the planning time.

In the pre-processing phase, prior to actual planning, the system uses the OOM in order to acquire all semantically relevant concepts for both the facts of the initial state and the outputs of the operators, discovered by the semantic analysis process described in the previous subsection. The enhancement of the problem by PORSCCE II is based on the following rules:

- The original concepts of the initial state together with the semantically equivalent and similar concepts form a new set of facts noted as the Expanded Initial State (EIS).
- The goals of the problem remain the same.
- The Enhanced Operator Set (EOS) is produced, by altering the description of each operator, while preserving the initial size of the set. More specifically, the preconditions of each operator remain the same, while the list of add effects of each operator is enhanced by including all the equivalent and semantically similar concepts for the concepts in the initial list.

Suppose, for example, that the initial state I and the two operators of the problem are the following:

```
I = {debitcard(X), dates, motel}
ActivateCard: prec={creditcard(X), disabled(X)},
              effects(+)={enabled(X)}, effects(-)={disabled(X)}
BookHotel: prec={dates, hotel}, effects(+)={bookinginfo}, effects(-)={},
```

The OOM for a given distance metric and threshold discovers the following relevant concepts:

```
debitcard ~ creditcard    motel ~ hotel    active ~ enabled
```

The pre-processor alters the problem definition to the following:


```

EIS: {debitCard(X), dates, motel, creditCard(X), hotel}
EOS: ActivateCard: prec={creditCard(X), disabled(X)},
      effects(+)={enabled(X), active(X)}, effects(-)={disabled(X)}
      BookHotel: prec={dates, hotel}, effects(+)={bookinginfo}, effects(-)={}
    
```

The new problem, namely $\langle \text{EIS}, \text{EOS}, \text{G} \rangle$ is encoded into PDDL and forwarded to the planning system in order to acquire a solution. Note that the semantic information is encoded in such a way that it is transparent to the external planners, which can solve the problem as any other classical planning problem.

5.4 PDDL to OWL-S Translation

After the acquisition of solutions, a reverse translation process has to take place, in order to provide the resulting composite web service to the original OWL-S standard and the initial web services domain. This reverse translation accommodates composite service deployment and execution monitoring.

For the purposes of the PORSCE II framework, the use of intricate OWL-S control constructs is not mandatory, as far as the proper invocation of the atomic processes is concerned. Since the modeling of the web service composition problem to a planning problem is merely based on the STRIPS formalism, there is no need to define alternative paths. Moreover, all the plans produced by the planning systems contain a finite number of steps and the use of loops is rare and not mandatory. Therefore, any plan produced by the framework can be expressed as a composite web service by using only the Sequence control construct, without risking the proper invocation of the composite service. This is true even for the cases where the external planning system used, such as *LPG-td*, returns a non linear plan (i.e. one that contains steps with parallel execution of actions), since any non linear STRIPS plan has one or more equivalent topological orderings. However, in order to accelerate the invocation of the composite service by allowing the parallel execution of certain atomic processes, an algorithm that translates plans (linear or non linear) to composite web services using the Sequence, Split and Split+Join constructs has been developed.

Algorithm 1 (Basic) Computes an initial composite service with *Sequence* and *Split* constructs
Inputs: $G = (V, E)$, the web service graph
Output: C : a composite service

```

1  set  $R \leftarrow \{r \in V : \forall x \in V, (x \rightarrow r) \notin E\}$  //  $R$  is the set of root nodes in  $G$ 
2  if  $|R| = 0$  then return NULL
3  if  $|R| = 1$  then
    set  $G' \leftarrow$  the tree in  $G$  with  $r \in R$  as the root
    return sequence( $r$ , Basic( $G' - \{r\}$ ))
4  set  $c \leftarrow \{\}$ 
5  for each  $r$  in  $R$ 
    set  $G' \leftarrow$  the tree in  $G$  with  $r \in R$  as the root
    set  $c \leftarrow c \cup$  Basic( $G' - \{r\}$ )
6  return split( $c$ )
    
```

Algorithm 1 presents the basic algorithm that creates a composite service, given a web service graph. A web service graph is a graph $G = (V, E)$, where the nodes in V correspond to all the atomic services in the plan and the edges $(x \rightarrow y)$ in E , where x and y are nodes in V , define that web service x produces an output that is required by y as an input. The process of obtaining a web service graph from the plan is straightforward and due to space limitations, we will not further elaborate on that. The output of the *Basic* function in Algorithm 1 is either a composite construct of the form *sequence*(c_1, c_2), or *split*(c_1, c_2, \dots, c_n), where c_1 to c_n are either *NULL* or composite constructs. For example, consider the web service graph presented in Fig. 4. The output of the *Basic* function presented in Algorithm 1 will be *split*(*sequence*($WS_a, \text{sequence}(WS_c, \text{NULL})$), *sequence*($WS_b, \text{sequence}(WS_c, \text{NULL})$))




Figure 4 Web service graph example.

Algorithm 2 (Join) Replaces *Split* with *Split+Join* where possible in a composite service

Inputs: $C = f(a_1, a_2, \dots, a_n)$: a composite service with *Sequence* and *Split* constructs

Output: C' : a composite service with *Sequence*, *Split* and *Split+Join* constructs

```

1  set  $f(a_1, a_2, \dots, a_n) = C$ , where  $f$  is the name of the construct and  $a_1$  to  $a_n$  its arguments
2  if  $f = NULL$  then return  $NULL$ 
3  if  $f = sequence$  then
     $a'_1 = Join(a_1)$ 
     $a'_2 = Join(a_2)$ 
    return  $f(a'_1, a'_2)$ 
4  if  $f = split$  then
    for each pair  $(a_i, a_j)$ ,  $i, j$  in  $[1, n]$ 
      if  $a_i$  and  $a_j$  have a common ending, i.e.  $a_i = a'_i \cup k$  and  $a_j = a'_j \cup k$  then
         $C' = C - \{a_i, a_j\} \cup seq(split + join(a'_i, a'_j), k)$ 
    return  $C'$ 

```

The output of the *Basic* function is then fed to the *Join* function, presented in Algorithm 2, in order to replace the *Split* construct with *Split+Join* wherever this is possible. The *Join* function searches in all possible pairs of *Split* arguments, in order to find a common ending part. For instance, in the example composite service given above, both arguments of the *Split* construct end in $sequence(WS_c, NULL)$. Therefore, if we apply Algorithm 2 to the output of Algorithm 1, the resulting composite service will be $split(sequence(split + join(WS_a, WS_b), sequence(WS_c, NULL)))$

The above composite service is then simplified by removing *NULLs* and constructs with single arguments and the final outcome is a construct of the form $sequence(split + join(WS_a, WS_b), WS_c)$

The above algorithms, along with some additional filters, such as the one mentioned above for removing *NULLs* and unary constructs were implemented using the CMU OWL-S API (CMU).

6 Solution and Integration

PORSCE II aims at a high degree of integration of the composition process; therefore, its features include solving the problem through invocation of external planning systems, visualization, solution evaluation and composite service modification.

6.1 Acquiring Solutions

Since the transformation process results in the export of both the planning domain and problem in PDDL, any PDDL-compliant domain independent external planning system can be used. This is a key issue for the ability of the system to keep up to date with advancements in planning research. Currently, two different planning modules have been incorporated in the system: JPlan (JPlan), which is an open-source Java implementation of Graphplan and LPG-*td* (Gerevini et al., 2004). Both planners proved to be remarkably fast and can handle a respectable number of operators, which is very important as the number of available web services is expected to increase significantly over time. After the planning process is completed, JPlan provides the plan, in its own format, which comprises of a simple sequential list of actions. LPG-*td*, on the other hand, provides the plan in a format that complies with PDDL+. The plan in this case might not be sequential, but structured in levels; actions belonging to the same level can be executed in an arbitrary sequence, however all actions of a certain level must be completed before any action of the following level can be executed. Subsequently, the produced plans are visualized and their accuracy is evaluated.

6.2 Composite Service Accuracy Assessment

Semantic relaxation and the use of multiple planners may produce a number of composite services, for which statistics and quality metrics have to be calculated. Such metrics include the number of actions and the number of levels in the plan, as well as a plan distance quality metric, which indicates the accuracy of the plan, when semantic relaxation takes place.

For the calculation of the plan semantic distance, each concept appearing in the inputs or outputs of the actions of the plan is annotated by the OOM with a semantic distance d_i with respect to the original

concept it was derived from, using the selected similarity metric. A concept distance of 0 reveals identical or equivalent concepts. Additionally, each concept is annotated with a weight w_i , with respect to the kind of hierarchical relationship to the original concept. This provides the option to discriminate among different hierarchical relationships, to accommodate for cases when certain relationships might be more desirable than others. As a rough example consider a user looking for "zip_code" (subclass); in this case, superclasses, for example "address" (which entails the zip code) are more desirable than siblings, for example "street_number"; therefore, providing different weights for superclasses and siblings will have the desired effect on the results. These values are combined to form the plan semantic distance. When the upwards cotopic distance metric is used, the plan semantic distance is calculated as a weighted product of these concepts, as the product represents appropriately the semantic distance in this case:

$$PSD_{uc} = \prod_{k=0}^n w_i d_i, d_i \neq 0$$

The plan accuracy metric is calculated as $1 - PSD$; therefore, if there is exact input to output matching, or if only equivalent concepts are used, then the plan quality metric value is 1, while it decreases as the plan becomes less accurate.

6.3 Visualization and Composite Service Modification

The Visualizer enhances comprehensibility by providing a visual representation of the composite service and enables the user to interfere by manipulating it. The composite service is represented as a schema of simple service invocations, possibly structured in levels, showing inputs and outputs, as well as dependencies among web services. The Visualizer module invokes and interacts closely with another module of the system, the Service Replacement Component, which allows the user to select among a series of alternatives in order to modify the produced composite service. The first alternative for composite service modification is the replacement of a certain service included in the composite service (plan) with a semantically equivalent or relevant service. In order to perform this operation, the system needs to discover all actions that could be used alternatively instead of the chosen one, using advice from the OOM as far as concept equivalence and semantic relevance are concerned. An action A is considered an alternative for an action Q of the plan as far as it does not disturb the plan sequence and the intermediate states. In order to ensure that, both the conditions $add(A) \supseteq add(Q)$ and $prec(A) \subseteq S$ must hold, where S is the state of the world exactly before the application of A and can be calculated by Algorithm 3 below (I_{rr}). The selected alternative service substitutes the original one both in the plan and in the visualization, and no replanning is performed. The web service substitution can be applied an arbitrary number of times on any of the services taking part in the composite service. In cases when none of the semantically equivalent or relevant services that correspond to a certain service is considered suitable, or in cases where there are no alternative services, the system offers the option to substitute the service with a partially ordered set of services, which are found through planning. In this case, the world states right before and after the execution of the action being replaced serve as the initial and goal states for the planning process, respectively. In order to find the initial state I_{rr} and the goal state G_{rr} for the replanning process Algorithms 3 and 4 are used respectively. Note that the replanning process is bound to return the web service being replaced itself, especially if the external planner used produces the optimal plan in each case. In order to prevent that, this specific service has to be removed from the set of available services before the replanning process proceeds. As the new plan produced substitutes the service, its quality metrics have to be incorporated in the quality metrics of the entire plan.

If replacement of a web service, either by an equivalent or through replanning, is not a suitable option, or if multiple services are undesirable or unavailable, the user can resort to replanning from a certain point in the plan, or even replanning from scratch. When replanning from a certain point, Algorithm 3 is used to calculate the initial state.

7 Demonstration and System Evaluation

This section aims at demonstrating the use and evaluating the performance of the PORSCE II through a case study, following the general course depicted in Fig. 5.

Algorithm 3 Computes the initial state for replacement of action A_i through replanning (I_{rr})	
Inputs: Extended Initial State (EIS), plan $P = A_1, \dots, A_n$, index of the action being replaced (i)	
Output: The I_{rr}	
1	set $I_{rr} \leftarrow EIS$
2	if $j = 1$ // start at the beginning of the plan
3	do // for every action
	$I_{rr} \leftarrow I_{rr} \cup add(A_j)$ // include the add effects of the action in the I_{rr} set
	$j = j + 1$
4	while $j < i$ // until the action being replaced is reached
5	return I_{rr}
<hr/>	
Algorithm 4 Computes the goal state for replacement of action A_i through replanning (G_{rr})	
Inputs: Initial Goal State (G), plan $P = A_1, \dots, A_n$, index of the action being replaced (i)	
Output: The G_{rr}	
1	set $G_{rr} \leftarrow G$
2	if $j = n$ // start at the end of the plan
3	do // for every action
	$G_{rr} \leftarrow G_{rr} - add(A_j) \cup prec(A_j)$ // include preconditions and exclude the add effects
	$j = j - 1$
4	while $j > i$ // until the action being replaced is reached
5	return G_{rr}

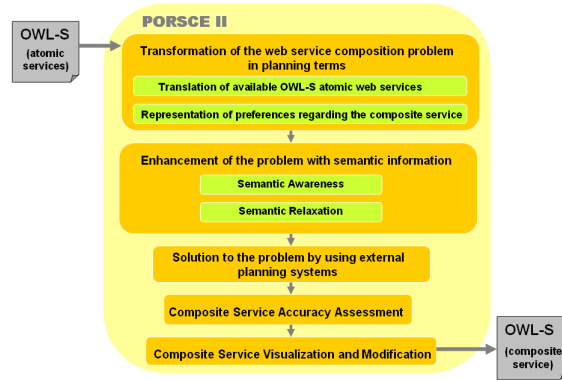


Figure 5 The demonstration steps.

The test sets used to perform experiments were obtained from the OWLS-TC (OWLS-TC). Several service descriptions were modified or added to the domains, accommodating demonstration of the capabilities of the system. Some indicative web services that were modified or added are depicted in Table 2.

Service	Inputs	Preconditions	Outputs	Effects
BookToPublisher	Book, Author		Publisher	
CreditCardCharge	OrderData, CreditCard		Payment	
ElectronicOrder	Electronic		OrderData	
PublisherElectronic Order	PublisherInfo		OrderData	
ElectronicOrderInfo	Electronic		OrderInformation	
Shipping	Address, OrderData		ShippingDate	
WaysOfOrder	Publisher		Electronic	
CustomsCost	Publisher, OrderData		CustomsCost	
FindUsed	ISBN		Seller	has(Seller,ISBN)
OrderUsed	ISBN, Seller, Client	has(Seller,ISBN)	OrderData	-has(Seller,ISBN) has(Client,ISBN)
FindISBN	Book, Author		ISBN	

Table 2 Added / modified web services.

The transformation of the web service composition problem in planning terms includes translating all available OWL-S web services, including the aforementioned ones, to PDDL operators, so that the size of the resulting domain is maintained on realistic levels. The transformation process also incorporates the representation of the requirements about the composite service, which the user can express through a dialog interface such as the one depicted in Fig. 6.

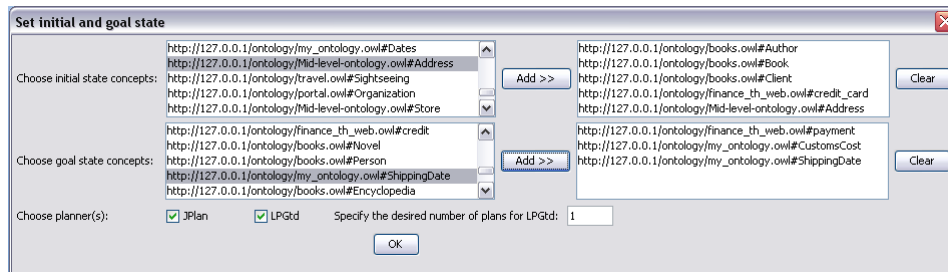


Figure 6 Initial and goal states definition and desired planners selection.

The scenario implemented here belongs to the OWLS-TC *books* and *finance* domains, and concerns the electronic purchase of a book. The user provides as inputs his details (client), a book title and author, credit card info and the address that the book will be shipped to, and wishes to use a credit card for the purchase, as well as to be informed about the shipping dates and the customs cost for the specific item. The initial state corresponds to the inputs of the composite service, while the goal state represents the desired composite service outcome.

In order to accommodate semantic awareness, all the ontologies that organize the concepts appearing as inputs and outputs of the available web services are parsed and analyzed. This enables semantic relaxation, performed through semantic enhancement of the planning domain and problem. The degree of the semantic relaxation is user-defined, and can be specified by selecting semantic distance metrics and thresholds through the interface depicted in Fig. 7.

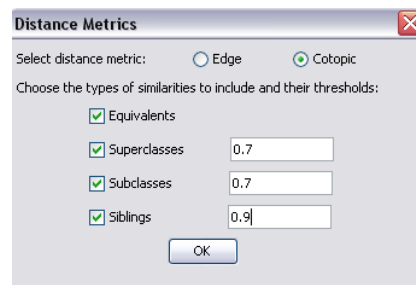


Figure 7 The semantic enhancement interface.

At this point, the system exports the formulated and possibly semantically enhanced planning domain and problem to PDDL. Consequently, it invokes external planners to acquire solutions. The PDDL domain, problem and produced plan for this scenario are depicted in Fig. 8. Note that the domain in this case, for space purposes, contains only the necessary atomic web services.

The produced plans are imported into the Visualizer Component, where they are represented as a web service graph and depicted visually. The first plans produced by JPlan and LPG-td for this case study, using the operator set described above, without performing any semantic relaxation, are presented in Fig. 9. The calculated statistics and metrics for the composite web services include the number of actions and the number of levels in the plan (which coincide for sequential plans), as well as the plan accuracy metric.

While exact matching of input to output concepts is obligatory in the classical planning domains, in the web services world the case can be different, as it is preferable to present the user with a composite service that approximates the required functionality than to present no service at all. The semantically similar concepts obtained from the OWL Ontology Manager enable the system to compose alternative services that approximate the desired one in case there are no exact matches, by performing semantically

<pre> DOMAIN: (define (domain wscomp) (:predicates (block ?b) (ontologybooksowlAuthor ?b) (ontologybooksowlBook ?b) (ontologybooksowlPublisher ?b) (ontologyfinancethwebowlcreditcard ?b) (ontologymyontologyowlOrderData ?b) (ontologyfinancethwebowlpayment ?b) (ontologymyontologyowlCustomsCost ?b) (ontologyfinancethwebowlElectronic ?b) (ontologyMidlevelontologyowlAddress ?b) (ontologymyontologyowlShippingDate ?b) (ontologybooksowlISBN ?b) (ontologybooksowlSeller ?b) (ontologybooksowlClient ?b) (has_Seller_ISBN ?b) (has_Client_ISBN ?b)) (:action BookToPublisherService :parameters (?b) :precondition (and (block ?b) (ontologybooksowlAuthor ?b) (ontologybooksowlBook ?b)) :effect (ontologybooksowlPublisher ?b)) (:action CreditCardChargeService :parameters (?b) :precondition (and (block ?b) (ontologyfinancethwebowlcreditcard ?b) (ontologymyontologyowlOrderData ?b)) :effect (ontologyfinancethwebowlpayment ?b)) (:action CustomsCostService :parameters (?b) :precondition (and (block ?b) (ontologymyontologyowlOrderData ?b)) :effect (ontologymyontologyowlCustomsCost ?b)) (:action ElectronicOrderService :parameters (?b) :precondition (and (block ?b) (ontologyfinancethwebowlElectronic ?b)) :effect (ontologymyontologyowlOrderData ?b)) (:action ShippingService :parameters (?b) :precondition (and (block ?b) (ontologyMidlevelontologyowlAddress ?b) (ontologymyontologyowlOrderData ?b)) :effect (ontologymyontologyowlShippingDate ?b)) (:action WaysOfOrderService :parameters (?b) :precondition (and (block ?b) (ontologybooksowlPublisher ?b)) :effect (ontologyfinancethwebowlElectronic ?b)) (:action FindUsed :parameters (?b) :precondition (and (block ?b) (ontologybooksowlISBN ?b)) :effect (and (ontologybooksowlSeller ?b) (has_Seller_ISBN ?b))) (:action FindISBN :parameters (?b) :precondition (and (block ?b) (ontologybooksowlBook ?b) (ontologybooksowlAuthor ?b)) :effect (ontologybooksowlISBN ?b)) (:action OrderUsed :parameters (?b) :precondition (and (block ?b) (ontologybooksowlISBN ?b) (ontologybooksowlSeller ?b) (ontologybooksowlClient ?b) (has_Seller_ISBN ?b)) :effect (and (ontologymyontologyowlOrderData ?b) (has_Client_ISBN ?b) (not (has_Client_ISBN ?b)))) </pre>	<pre> PROBLEM: (define (problem wscompprob) (:domain wscomp) (:objects a) (:init (block a) (ontologybooksowlAuthor a) (ontologybooksowlBook a) (ontologyfinancethwebowlcreditcard a) (ontologybooksowlClient a) (ontologyMidlevelontologyowlAddress a)) (:goal (and (ontologyfinancethwebowlpayment a) (ontologymyontologyowlCustomsCost a) (ontologymyontologyowlShippingDate a)))) </pre>
	<pre> PLANS (LPG-td) : ; Version LPG-td-1.0 ; Seed 52716286 ; Command line: lpg-td-1.0 -o domain.pddl -f problem.pddl -n 1 ; Problem problem.pddl ; Actions having STRIPS duration ; Time 0.01 ; Search time 0.00 ; Parsing time 0.01 ; Mutex time 0.00 ; Quality 6 Time 0.01 0: (BOOKTOPUBLISHERSERVICE A) [1] 1: (WAYSOFORDERSERVICE A) [1] 2: (ELECTRONICORDERSERVICE A) [1] 3: (CREDITCARDCHARGESERVICE A) [1] 3: (CUSTOMSCOSTSERVICE A) [1] 3: (SHIPPINGSERVICE A) [1] OR ; Version LPG-td-1.0 ; Seed 46626426 ; Command line: lpg-td-1.0 -o domain.pddl -f problem.pddl -n 1 ; Problem problem.pddl ; Actions having STRIPS duration ; Time 0.03 ; Search time 0.02 ; Parsing time 0.02 ; Mutex time 0.00 ; Quality 6 Time 0.03 0: (FINDISBN A) [1] 1: (FINDUSED A) [1] 2: (ORDERUSED A) [1] 3: (CREDITCARDCHARGESERVICE A) [1] 3: (CUSTOMSCOSTSERVICE A) [1] 3: (SHIPPINGSERVICE A) [1] </pre>

Figure 8 The PDDL domain, problem and plan for the specific scenario.

relaxed concept matching. Such an approximate service for the specific case study is presented in Fig. 10. The calculated accuracy of this service is different from the accurate ones presented in Fig. 9.

In case service replacement is required, for example on the CustomsCost service, and there is no alternative service available, service replacement through replanning will be employed. The algorithms described in the corresponding section will yield new initial and goal states, and the corresponding planner will be re-invoked, finding a new sequence of actions that can substitute the selected service. The user interface for the replacement options is depicted in Fig. 11, while the resulting composite service after the modifications is depicted in Fig. 12.

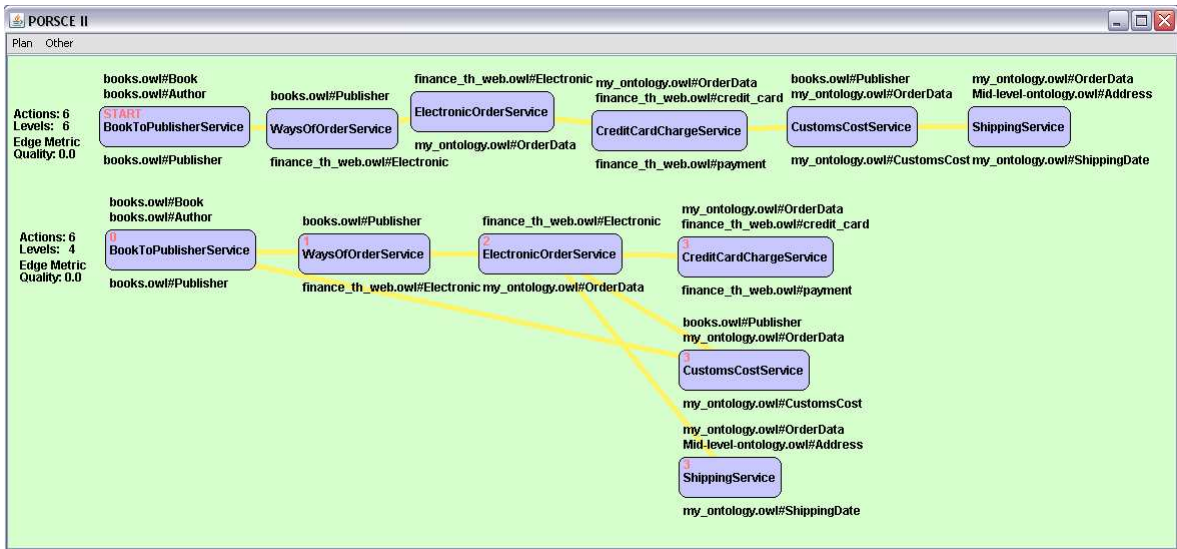


Figure 9 The plans from JPlan (top) and LPG-td (bottom) for the specific case study.

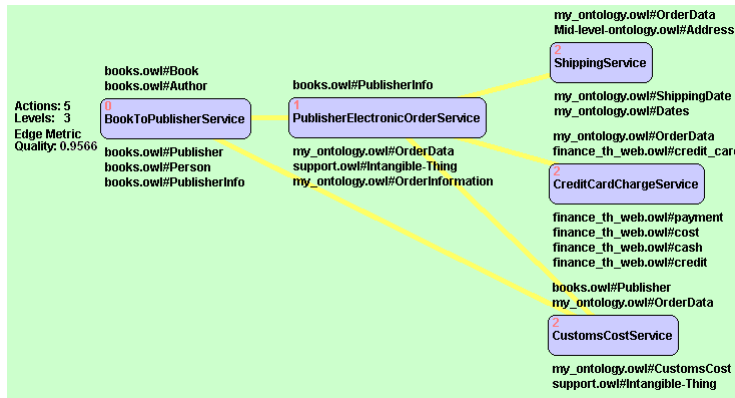


Figure 10 Approximate composite service.




Figure 11 Service Replacement Interface.




Figure 12 Composite service after replacement operation.

The final step is to transform the solution to web service context. This is achieved by translating the PDDL+ plan that represents the desired composite web service, produced by the external planning systems, into OWL-S, utilizing information retrieved from web service descriptions and ontology analysis.

Number of web services		10	100	500	1000
Preprocessing time		5953	5794	6007	5903
Total transformation time	Exact	4685	71533	356800	823793
	Edge	4635	76558	346212	820700
	Cotopic	4667	76928	757778	3947955
Transformation time per web service	Exact	469	715	714	824
	Edge	464	766	693	820
	Cotopic	467	769	1516	3950
Planning time (LPG-td)	Exact	4	17	50	119
	Edge	6	16	72	123
	Cotopic	4	15	56	122

Table 3 Time measurements in milliseconds.

In order to study the behavior of the system as the number of available web services increases, web service profiles were added to the domain progressively in batches. The time performance results presented in Table 3 were obtained from a number of runs of the system on a machine with Dual-Core AMD Opteron Processor at 2.20GHz with 1GB of RAM memory and concern times for preprocessing, transformation of the OWL-S service profiles to PDDL actions and planning using *LPG-td*.

Measurements took place for domains of different sizes, namely 10, 100, 500 and 1000 OWL-S profiles. Some of the experiments were performed with exact matching, while others were performed with semantic relaxation using either the edge-counting or the upwards cotopic metric. The preprocessing time did not show significant fluctuation, as it depends on the number and structure of the processed ontologies and not on the number of available web services. The total transformation time evidently increased as the number of available web services increased, however the average transformation time per web service profile converged to approximately 0.81 seconds for the exact matching and the edge-counting cases. In the upwards cotopic case, the increase in the average transformation time is significant as available web services increase, due to the higher complexity of the algorithm used for calculating the upwards cotopic relevance between two concepts. As far as average planning time is concerned, *LPG-td* shows an increase in planning time as the number of actions increases; however, it is still remarkably fast. It should be noted that the scalability of the system as far as planning time is concerned is not critical, as the planners used are not embedded and could easily be replaced by more efficient ones.

8 Conclusions and Future Work

This paper presented *PORSCE II*, an integrated system which combines planning with semantic object relevance in order to approach automated semantic web service composition. The web service composition problem is transformed into a planning problem, solved under semantic awareness accommodating approximate compositions, and then transformed back in web service terms. The system exploits the most prominent standards in both worlds, namely OWL-S and PDDL. *PORSCE II* aims at a high degree of interoperability with external planning systems which perform planning with the desired degree of semantic relaxation. Finally, the system is integrated with a visual environment and components which accommodate composite service evaluation and modification. Among the main advantages of the proposed framework are the extended utilization of semantic information, the ability to scale up for a great number of services, and the capability to handle service failure or unavailability dynamically.

Future goals include the extension of the system in order to deploy the produced composite services, through OWL-S deployment systems such the OWL-S Virtual Machine (Paolucci et al., 2003), and automatically acquire feedback, which can then be utilized to partially automate the service replacement procedure. Also, another direction we plan to explore is the modification of the web service modeling, so that a metric domain is constructed, which incorporates the concept semantic distances and requires the planner to minimize the total semantic distance. In addition, another future goal concerns the exploration of the possibility to accelerate the composition process by asserting the produced OWL-S profiles in the base of the available web services, under certain time constraints. Furthermore, integration with the *VLEPPO* system (Hatzi et al., 2007) is a promising future direction, in order to accommodate design and solving of the web service composition problems. Finally, it lies in our immediate plans to study

ways to enhance the services representation and explore the ability to produce various composite services according to non-functional properties.

References

- CMU. OWL-S API. <http://www.daml.ri.cmu.edu/owlapi/> [accessed 17/04/10].
- S. Dustdar and W. Schreiner. A survey on web services composition. *Int. J. Web Grid Services*, 1(1):1–30, 2005.
- R. E. Fikes and N. J. Nilsson. STRIPS: a new approach to the application of theorem proving to problem solving. In *IJCAI'71: Proceedings of the 2nd international joint conference on Artificial Intelligence*, pages 608–620, San Francisco, CA, USA, 1971. Morgan Kaufmann Publishers Inc.
- M. Fox and D. Long. PDDL+: Modelling continuous time-dependent effects, 2002.
- A. Gerevini and D. Long. Plan constraints and preferences in PDDL3. Technical report, Department of Electronics for Automation, University of Brescia, Italy, 2005.
- A. Gerevini, A. Saetti, I. Serina, and P. Toninelli. LPG-td: a fully automated planner for pddl2.2 domains. In *In Proc. of the 14th Int. Conference on Automated Planning and Scheduling (ICAPS-04) International Planning Competition abstracts*, 2004.
- M. Ghallab, A. Howe, C. Knoblock, D. McDermott, A. Ram, M. Veloso, D. Weld, and D. Wilkins. PDDL – the planning domain definition language. Technical report, Yale University, New Haven, CT, 1998.
- O. Hatzi, D. Vrakas, N. Bassiliades, D. Anagnostopoulos, and I. Vlahavas. VLEPPO: A visual language for problem representation. In R. Bartak, editor, *PlanSIG 2007: The 26th workshop of the UK Planning and Scheduling Special Interest Group*, pages 60–66, 2007.
- O. Hatzi, G. Meditskos, D. Vrakas, N. Bassiliades, D. Anagnostopoulos, and I. Vlahavas. A synergy of planning and ontology concept ranking for semantic web service composition. In *IBERAMIA '08: Proceedings of the 11th Ibero-American conference on AI*, pages 42–51, Berlin, Heidelberg, 2008. Springer-Verlag.
- O. Hatzi, G. Meditskos, D. Vrakas, N. Bassiliades, D. Anagnostopoulos, and I. Vlahavas. PORSCE II: Using planning for semantic web service composition. In R. Bartak, S. Fratini, and L. McCluskey, editors, *In Proceedings of 3rd International Competition on Knowledge Engineering for Planning and Scheduling (ICKEPS'09), in conjunction with the International Conference on Automated Planning and Scheduling (ICAPS-09)*, pages 38–45, 2009.
- IPC2004. International Planning Competition. <http://www.tzi.de/edelkamp/ipc-4/> [accessed 27/10/10].
- JPlan. Java Graphplan Implementation. <http://sourceforge.net/projects/jplan> [accessed 27/10/10].
- M. Klusch and A. Gerber. Semantic web service composition planning with OWLS-XPlan, 2005.
- A. Maedche and V. Zacharias. Clustering ontology-based metadata in the semantic web. In *PKDD '02: Proceedings of the 6th European Conference on Principles of Data Mining and Knowledge Discovery*, pages 348–360, London, UK, 2002. Springer-Verlag.
- D. McDermott. Estimated-regression planning for interactions with web services. In *In Proceedings of the 6th International Conference on Artificial Intelligence Planning Systems*, pages 204–211. AAAI Press, 2002.
- S. McIlraith and T. C. Son. Adapting GOLOG for composition of semantic web services. In *Proceedings of the Eighth International Conference on Knowledge Representation and Reasoning*, pages 482–493, 2002.
- OWL. Web Ontology Language. <http://www.w3.org/TR/owl-ref/> [accessed 27/10/10].
- OWL-S. 1.1. <http://www.daml.org/services/owl-s/1.1/> [accessed 27/10/10].
- OWLS-TC. SemWebCentral. <http://projects.semwebcentral.org/projects/owl-s-tc/> [accessed 27/10/10].
- M. Paolucci, A. Ankolekar, N. Srinivasan, and K. Sycara. The DAML-S virtual machine. In *The SemanticWeb - ISWC 2003*, volume 2870 of *Lecture Notes in Computer Science*, pages 290–305. Springer Berlin / Heidelberg, 2003.
- M. Pistore, A. Marconi, P. Bertoli, and P. Traverso. Automated composition of web services by planning at the knowledge level. In *In 19th Intl. Joint Conferences on Artificial Intelligence*, pages 1252–1259, 2005.
- S. R. Ponnekanti and A. Fox. SWORD: A developer toolkit for web service composition. In *Proceedings of the 11th International WWW Conference (WWW2002)*, pages 83–107, Honolulu, HI, USA, 2002. Elsevier.
- J. Rao and X. Su. A survey of automated web service composition methods. In *In Proceedings of the 1st International Workshop on Semantic Web Services and Web Process Composition, SWSWPC 2004*, pages 43–54, 2004.
- RuleML. The Rule Markup Initiative. <http://ruleml.org/> [accessed 27/10/10].
- SAWSDL. Semantic Annotations for WSDL. <http://www.w3.org/2002/ws/sawSDL/> [accessed 27/10/10].
- E. Sirin, B. Parsia, D. Wu, J. Hendler, and D. Nau. HTN planning for web service composition using SHOP2. *Web Semantics*, 1(4):377–396, 2004.
- E. Sirin, B. Parsia, B. C. Grau, A. Kalyanpur, and Y. Katz. Pellet: A practical OWL-DL reasoner. *Journal of Web Semantics*, 5(2):51–53, 2007.
- SWRL. A Semantic Web Rule Language. <http://www.w3.org/Submission/SWRL/> [accessed 27/10/10].