

Chapter 4

Semantic Web Service Coordination

Matthias Klusch

4.1 Introduction

Semantic service coordination aims at the coherent and efficient discovery, composition, negotiation, and execution of Semantic Web Services in a given environment and application context. What makes coordination of services in the Semantic Web different from its counterpart in the Web is its far more advanced degree of automation through means of logic-based reasoning on heterogeneous service and data semantics.

In this chapter, we only focus on approaches to semantic discovery and composition planning of Semantic Web services, and briefly comment on their interrelationships and selected open problems of both fields. For reasons of space limitations, the set of presented examples is representative but not exhaustive.

4.2 Semantic Service Discovery

Service discovery is the process of locating existing Web services based on the description of their functional and non-functional semantics. Discovery scenarios typically occur when one is trying to reuse an existing piece of functionality (represented as a Web service) in building new or enhanced business processes. A Semantic Web service, or in short semantic service, is a Web service which functionality is described by use of logic-based semantic annotation over a well-defined ontology (cf. Chapter 3). In the following, we focus on the discovery of semantic services. Both service-oriented computing and the Semantic Web envision intelligent agents to proactively pursue this task on behalf of their clients.

Semantic service discovery can be performed in different ways depending on the considered service description language, means of service selection and

coordination through assisted mediation or performed in a peer-to-peer fashion. In general, any service discovery framework needs to have the following components ([37]).

- **Service description language:** A service description language (more precisely top-level ontologies, also called service description formats) is used to represent the functional and non-functional semantics of Web services. Examples of structured and logic-based semantic service description language are OWL-S and WSML. The standard Semantic Web service description language SAWSDL allows for a structured representation of service semantics in XML(S) with references to any kind of non-logic-based or logic-based ontology for semantic annotation.¹ Alternatively, in so-called monolithic logic-based service descriptions the functionality of a service is represented by means of a single logical expression of an appropriate logic, usually a description logic like OWL-DL or WSML-DL.
- **Service selection means:** Service selection encompasses semantic matching and ranking of services to select a single most relevant service to be invoked, starting from a given set of available services. This set can be collected and maintained, for example, by front-end search engine, or given by providers advertising their services at registries or middle-agents like matchmakers and brokers. Semantic service matching, or in short: service matching, is the pairwise comparison of an advertised service with a desired service (query) to determine the degree of their semantic correspondence (semantic match). This process can be non-logic-based, logic-based or hybrid depending on the nature of reasoning means used.

Non-logic-based matching can be performed by means of, for example, graph matching, data mining, linguistics, or content-based information retrieval to exploit semantics that are either commonly shared (in XML namespaces), or implicit in patterns or relative frequencies of terms in service descriptions. Logic-based semantic matching of services like those written in the prominent service description languages OWL-S (Ontology Web Language for Services), WSML (Web Service Modeling Language) and the standard SAWSDL (Semantically Annotated WSDL) exploit standard logic inferences. Hybrid matching refers to the combined use of both types of matching.

- **Discovery architecture:** The conceptual service discovery architecture concerns the environment in which the discovery is assumed to be performed. This includes assumptions about the (centralized or decentralized P2P) physical or semantic overlay of the network, the kind of service information storage (e.g., service distribution, registries, and ontologies) and location mechanisms such as query routing, as well as the agent society in the network (e.g., service consumers, providers, middle-agents).

¹In this sense, SAWSDL services can be seen as a weaker form of semantic services while WSDL services are no semantic services.




Figure 4.1: Categories of existing Semantic Web Service matchmakers.

In the following, we survey existing approaches to semantic service matching and discovery architectures. Examples of semantic service description languages were presented in the previous chapter.

4.2.1 Classification of Semantic Web Service Matchmakers

Semantic service matching determines whether the semantics of a desired service (or goal) conform to that of an advertised service. This is at the very core of any semantic service discovery framework. Current approaches to semantic service matching can be classified according to

- what kinds and parts of service semantics are considered for matching, and
- how matching is actually performed in terms of non-logic-based or logic-based reasoning on given service semantics or a hybrid combination of both, within or partly outside the respective service description framework (cf. Figure 4.1).

Non-Logic, Logic, and Hybrid Semantic Service Matching The majority of Semantic Web service matchmakers performs deductive, that is logic-based semantic service matching. In this sense, they are keeping with the original idea of the Semantic Web to determine semantic relations (thus resolve semantic heterogeneities) between resources including services based on logical inferencing on their semantic annotations that are formally grounded in description logics (DL) and/or rules (cf. Chapter 3). As shown in figure 4.1, pure logic-based semantic matchmakers for services in OWL-S and WSML are currently prevalent. Non-logic-based semantic service matchmakers do not perform any logic-based reasoning to determine the degree of a semantic match between a given pair of service descriptions. Examples of non-logic-based semantic matching techniques are text similarity measurement, structured graph matching, and path-length-based similarity of concepts².

Service Profile and Process Model Matching Most Semantic Web service matchmakers perform service profile rather than service process model matching. Service profile matching (so-called “black-box” service matching) determines the semantic correspondence between services based on the description of their profiles. The profile of a service describes what it actually does in terms of its signature, that is its input and output (IO), as well as preconditions (P) and effects or postconditions (E), and non-functional aspects such as the relevant business category, name, quality, privacy and pricing rules of the service. We classify additional context information for service matching such as the organisational (social or domain) roles, or geographic location of service requesters and providers in their interaction as non-functional.

Service process-oriented matching (so-called “glass-box” service matching) determines the extent to which the desired operational behavior of a given service in terms of its process control and data flow matches with that of another service. Like with service profile matching, we can distinguish between non-logic based, logic based and hybrid semantic process matching approaches depending on whether automated reasoning on operational semantics specified in some certain logic or process algebraic language (e.g. CCS, π -calculus) is performed, or not. An overview of relevant approaches to process mining for process discovery is given in [110].

Supported Semantic Web Service Description Formats Each of the implemented Semantic Web service matchmakers shown in Figure 4.1 supports only one of the many existing Semantic Web service description formats (cf. Chapter 3) as follows. This list is representative but not exhaustive.

²Please note that any kind of semantic service matching that identifies concepts or rules (which are logically defined in a given ontology) by their names only does not classify as logic-based matching in the strict sense. Without any formal verification of the semantic relation between given (semantic service annotation) concepts based on their logical definitions, the matchmaker performs non-logic-based semantic service matching.

- **OWL-S** matchmakers: Logic-based semantic matchmakers for OWL-S services are the OWLSM [45] and OWLS-UDDI [81] focussing on service IO-matching, and the PCEM [18] that converts given OWL-S services to PDDL actions for PROLOG-based service PE-matching. Further OWL-S matchmakers are the hybrid service IO-matchmaker OWLS-MX [59], the hybrid non-functional profile matchmaker ROWLS [34], the hybrid (combined) profile matchmaker FC-MATCH [16], the non-logic-based (full) service matchmaker iMatcher1 [14] and its hybrid successor iMatcher2 [51]. An approach to logic-based OWL-S process model verification is in [109] while [11] presents an approach to the matching of OWL-S process dependency graphs based on syntactic similarity measurements, and [12] proposes a hybrid matchmaker that recursively compares the DAML-S process model dependency graphs.
- **WSML** matchmakers: Implemented approaches to WSML service discovery include the hybrid semantic matchmaker WSMO-MX [49], the logic-based matchmaker GLUE [27], and the syntactic search engine for QoS-enabled WSML service discovery in P2P networks [112]. Other approaches to logic-based WSML service IOPE matchmaking are presented in [48, 103], though it is unclear to what extent they have been implemented.
- **WSDL-S/SAWSDL** matchmakers: The METEOR-S WSDI discovery infrastructure [111] and the UDDI-based search component Lumina³ are the only tool support of searching for SAWSDL services so far. While searching with Lumina is keyword-based, the MWSDI discovery of SAWSDL services relies on non-logic-based matching means.
- **Monolithic DL-based** matchmakers: Only very few matchmaker are agnostic to the above mentioned structured Semantic Web Service description formats without conversion by accepting monolithic descriptions of services in terms of a single service concept written in a given DL. In this case, semantic matching directly corresponds to DL inferencing, that is, semantic service matching is done exclusively within the logic theory such as performed by RACER [65], MaMaS⁴ [29, 30], and in [38]. Recently, an implemented approach to matching of monolithic service descriptions in OWL-DL extended with (non-functional) pricing policies modeled in DL-safe SWRL rules according to given preferences using SPARQL queries to a service repository is presented in [62].
- **Others**: Non-logic-based service IOPE profile matchmakers for other structured service description formats are the DSD matchmaker [53] for DIANE services, the numeric service IO type matching based HotBlu matchmaker [25], and the hybrid service IOPE matchmaker LARKS for services in an equally named format [105].

³lstdis.cs.uga.edu/projects/meteor-s/downloads/Lumina/

⁴sisinflab.poliba.it/MAMAS-tng/

In the following, we discuss each category of Semantic Web service matching together with selected representative examples of the above mentioned Semantic Web service matchmakers in more detail. This is complemented by a classification of existing service discovery architectures for which these matchmakers have been designed for, or can be used in principle. As stand-alone implementation, each matchmaker classifies as centralized service discovery system, though a few matchmaker have been also tested for, or were originally developed for decentralized P2P service retrieval systems like the OWLS-MX and the OWLS-UDDI matchmaker, respectively, the WSMO-QoS search engine and the DReggie/GSD matchmaker⁵.

4.2.2 Logic-Based Semantic Service Profile Matching

As mentioned above, logic-based semantic service matchmakers perform deductive reasoning on service semantics. The majority of such matchmakers pairwise compare logic-based descriptions of service profile semantics. In order to define these semantics, logical concepts and rules are taken from respective ontologies as first-order or rule-based background theories with a shared minimal vocabulary. Different ontologies of service providers and service requester are matched or aligned either at design time, or at runtime as part of the logic-based service matching process.

Matching Degrees

The degree of logic-based matching of a given pair of semantic service profiles can be determined either (a) exclusively within the considered logic theory by means of logic reasoning, or (b) by a combination of logical inferences within the theory and algorithmic processing outside the theory. Prominent logic-based matching degrees are exact, plugin, subsumes, and disjoint which are defined differently depending on the parts of service semantics and the logic theory that is used to compute these degrees.

One prominent example for a software specification matching degree is the so-called plug-in match. A specification S plugs into (plug-in matches with) another specification R , if the effect of S is more specific than that of R , and vice versa for the preconditions of S and R [115]. If this definition is restricted to effects only, the matching degree is called a post plug-in match. Unfortunately, the original notion of plug-in match has been adopted quite differently by most logic-based Semantic Web service matchmakers for both monolithic and structured service descriptions.

⁵For reasons of readability, the implemented (stand-alone) Semantic Web service matchmakers shown in Figure 4.1 each representing a central discovery system by itself are not again listed in Figure 4.2, and vice versa, that is, those matchmaking approaches being inherent part of the functionality of each node of decentralized discovery systems (but not available as stand-alone matchmaker) are not listed in Figure 4.1.

Monolithic Logic-Based Service Matching

Matching of monolithic logic-based semantic service descriptions (cf. Chapter 3) is performed exclusively by means of logic inferences within the considered logic theory. That is, the functionality of a Web service is represented by a single (monolithic) expression in an appropriate logic, usually a description logic like OWL-DL or WSML-DL. As a consequence, monolithic logic-based semantic service matching reduces to standard first-order (description) logic reasoning such as checking the satisfiability of service and query concept conjunction, or the entailment of concept subsumption over a given knowledge base. Furthermore, it is agnostic to any form of structured stateless (I/O) or stateful (IOPE) representation of service semantics like in OWL-S and WSML. The prominent degrees of semantic matching used by the majority of monolithic logic-based semantic service matchmakers are logic equivalence, post-plug-in match, subsumes, and fail.

For example, the logical so-called post-plug-in match of an advertised service S with a service request R bases on the entailment of concept subsumption of S by R over a given knowledge base kb extended by the axioms of S and R : $kb \cup S \cup R \models S \sqsubseteq R$. That is, the matchmaker checks if in each first-order interpretation (possible world) I of kb , the set S^I of concrete provider services (service instances) is contained in the set R^I of service instances acceptable to the requester: $S^I \subseteq R^I$. This assures the requester that each provided service instance offers at least the requested functionality, maybe even more. In other words, service S is more specific than the request R , hence considered semantically relevant. In contrast, the so-called logical subsumes match assures the requester that her acceptable service instances are also acceptable to the provider: $kb \cup S \cup R \models R \sqsubseteq S$.

Some monolithic DL-based service matchmakers also check for a so-called intersection or potential match (Grimm, 2007)[37]. This matching degree indicates the principled compatibility of service S with request R with respect to the considered knowledge base kb by means of either concept intersection or non-disjointness. In the first case, the advertised service concept S potentially matches with the (desired service or) query concept R if their concept conjunction $S \cap R$ is satisfiable with respect to kb in some possible world I such that $S^I \cap R^I \neq \emptyset$ holds. In the second case, the monolithic logic-based semantic service matchmaker makes a stronger check by determining whether the intersection of both concept extensions is non-empty in each possible world.

In general, the complexity of matching monolithic DL-based service descriptions is equal to the combined DL complexity. For example, post-plug-in matching of service concepts in OWL-Full, that is SHOIQ⁺ (including transitive non-primitive roles) has been shown to be undecidable [10] but decidable for OWL-DL, WSML-DL and DL-safe SWRL.

One problem of monolithic DL-based service matching is the risk to return false positives due to incomplete knowledge specified in service descriptions S, R or the domain ontology kb [38]. In other words, semantic matching of S with R with respect to kb based on monotonic DL reasoning under open-world assump-

tion (OWA) can wrongly succeed due to the existence of possible but unwanted interpretations of concepts or roles used in S or R over kb . Such unwanted possible worlds of kb are intuitively ruled out by humans by default - which accounts for their usually non-monotonic reasoning under closed-world assumption⁶.

One solution to this problem is to explicitly capture such default (common-sense) knowledge by adding, for example, appropriate concept disjointness axioms or object assertions to the knowledge base kb . This excludes possible worlds which are "obviously" wrong (but allowed due to open-world semantics) but is considered impracticable as it requires the modeler to somewhat "overspecify" the kb with "obvious" information.

An alternative solution is to perform semantic matching of services with local closed-world reasoning (cf. chapter 2) as proposed by [38]. Key idea is to exclude the unwanted possible worlds of knowledge base kb by means of an additional autoepistemic logic operator \mathbf{K} ⁷ that allows to restrict the interpretation of certain concepts C and roles r used in advertised and desired service descriptions S and R to named individuals (nominals) in the ABox of kb which are definitely known or not known to belong to them (C, r) ⁸. However, this local closure of concepts and roles in S, R for their interpretation in kb (i.e., locally closing off possible worlds of kb in S and R without any occurrence of \mathbf{K} in kb) under the local closed world-assumption (LCWA)⁹ by use of the \mathbf{K} -operator makes semantic matching dependent on the state of the world: It requires the existence of named individuals in the ABox of kb as representative (static) information on the locally closed con-

⁶The OWA states that the inability to deduce some fact from a knowledge base does not imply its contrary by default, that is, the fact may hold (not in all but) in some possible world (interpretations of kb). For example, the intersection match of $R = \textit{Flight} \sqcap \forall \textit{from}. \textit{UKCity}$ with $S = \textit{Flight} \sqcap \forall \textit{from}. \textit{USCity}$ with respect to the knowledge base $kb = \{\textit{UKCity} \sqsubseteq \textit{EUCity}, \textit{Flight} \sqsubseteq \exists \textit{from}. \top\}$ wrongly succeeds. The reason is that kb is underspecified in the sense that (due to the OWA) there can be possible worlds in which cities can be both in the UK and the US, which causes a false positive for the intersection match.

⁷The epistemic logic operator \mathbf{K} allows to refer to definitely known facts by intersecting all possible worlds: $(\mathbf{K}C)^{I,E} = \bigcap_{I \in E} C^{I,E}$. The epistemic concept $\mathbf{K}C$ is interpreted as the intersection of extensions of concept C over all first-order interpretations of kb , that is the set of all individuals that are known to belong to C (in the epistemic model $E(kb)$ of kb , that is the maximal non-empty set of all first-order interpretations of kb).

⁸In the above example, the intersection match of the request $R = \textit{Flight} \sqcap \forall \textit{from}. \mathbf{K} \textit{UKCity}$ with service $S = \textit{Flight} \sqcap \forall \textit{from}. \mathbf{K} \textit{USCity}$ with respect to the matchmaker knowledge base $kb = \{\textit{UKCity} \sqsubseteq \textit{EUCity}, \textit{Flight} \sqsubseteq \exists \textit{from}. \top, \textit{UKCity}(\textit{London})\}$ correctly fails, hence avoids to return a false positive. The satisfiability of the epistemic concept $S \sqcap R$ requires the existence of a named individual x in kb known to be both \textit{UKCity} and \textit{USCity} (that is kb entails $\textit{UKCity}(x) \sqcap \textit{USCity}(x)$, i.e. $kb \models \textit{UKCity}(x)$ and $kb \models \textit{USCity}(x)$, for every possible world I in the epistemic model $E(kb)$). While the named individual \textit{London} in the ABox of kb is definitely known to belong to the concept \textit{UKCity} , and also known to belong to \textit{EUCity} due to the inclusion axiom in the TBox of kb , it is not definitely known to also belong to \textit{USCity} ($kb \not\models \textit{USCity}(\textit{London})$). There is also no other named individual in kb which is both known to be in \textit{UKCity} and \textit{USCity} such that $S \sqcap R$ is not satisfied. An intersection match of R with different service $S' = \textit{Flight} \sqcap \forall \textit{from}. \mathbf{K} \textit{EUCity}$ correctly succeeds.

⁹The LCWA assumes that all individuals of some concept, or all pairs of individuals of some role are explicitly known in the local knowledge base (selected local concept or role closure).

cepts and roles¹⁰. Besides, using an autoepistemic extension of description logics like OWL-DL or WSML-DL for semantic service matching is still uncommon in practice, though (non-monotonic) reasoners such as for epistemic query answering in ALCK¹¹ can be easily integrated in a matchmaker.

Another application of non-monotonic reasoning to monolithic DL-based service matching is proposed in [24, 30]. The respective matchmaker MAMAS provides non-standard explanation services, that are non-monotonic logical abduction and contraction, for partial (also called approximated, intersection, or potential) matches. For example, concept contraction computes an explanation concept G to explain why a request concept R is not compatible with service concept S , that is, why $S \sqcap R$ is not satisfiable ($S \sqcap R \sqsubseteq \perp$). For this purpose, it keeps the least specific concept expression K of concept R such that K is still compatible with S , i.e. $\neg(K \sqcap S) \sqsubseteq \perp$. The remaining set G of constraints of R represents the desired explanation of mismatch. Such kind of non-monotonic logical service matching is NP-hard already for the simple description logic ALN. However, research in this direction has just begun and is, in part, related to research on non-monotonic reasoning with Semantic Web rule languages.

Examples of implemented monolithic DL-based matchmakers for service concepts written in OWL(-DL) and DAML+OIL are MAMAS [29, 30], respectively, RACER. Remarkably, both matchmakers determine the degree of post-plug-in match inverse to its original definition in [115].

Service Specification Matching

The logic-based semantic matching of service specifications (so-called PE-matching) concerns the comparison of their preconditions (P) and effects (E) and originates from the software engineering domain. As mentioned above, the plug-in matching of two software components S, R requires that the logic-based definition of the effect, or postcondition of S logically implies that of R , while the precondition of S shall be more general than that of R [115]. In other words, a logic-based semantic plug-in match of service specifications S, R requires (in every model of given knowledge base kb) the effect of advertised service S to be more specific than requested, and its precondition to be more general than requested in R . Depending on the Semantic Web service description framework (cf. Chapter 3), the logic language for defining service preconditions and effects ranges from, for example, decidable def-Horn (DLP), WSML-DL and OWL-DL to undecidable SWRL, KIF and F-Logic(LP).

¹⁰In the above example, the intersection match $S' \sqcap R$ would (wrongly) fail, hence causes a false negative, if the named individual *London* would not have been explicitly stated in kb to belong to *UKCity* as its representative by default: There would be no named individual definitely known to belong to both *UKCity* and *EUCity* in all possible worlds. Though *UKCity(London)* has to be added to the kb of the matchmaker to avoid false positives and negatives of intersection matches by non-monotonic epistemic query answering kb with **K**, no (dynamic) information about concrete flights, i.e. individuals of concept *Flight*, has to be additionally specified in kb .

¹¹<http://www.fzi.de/downloads/wim/KToy.zip>

For example, the logic-based service-PE matchmaker PCEM (cf. Chapter 10) exploits the Java-based light-weight Prolog engine tuProlog¹² for logic-based exact matching of service preconditions and effects written in Prolog. In particular, the PCEM matchmaker checks whether there is a possibly empty variable substitution such that, when applied to one or both of the logical propositions (PE), this results into two equal expressions, and applies domain specific inference rules (for computing subPartOf relations).

The hybrid semantic WSML service matchmaker WSMO-MX [49] is checking an approximated query containment over a given finite service instance base for service effects (postconditions, constraints) written in undecidable F-Logic(LP) using OntoBroker. The approach to semantic service IOPE matchmaking described in [103] uses the VAMPIRE theorem prover for matching pairs of preconditions and effects written in FOL, while the hybrid service IOPE matchmaker LARKS [105] performs polynomial theta-subsumption checking of preconditions and postconditions written in Horn. There are no non-logic-based or hybrid semantic service PE matchmaker available yet.

Service Signature and IOPE Matching

Logic-based semantic matching of service signatures (input/output, IO), so called service profile IO-matching, is the stateless matching of declarative data semantics of service input and output parameters by logical reasoning within the theory and algorithmic processing outside the theory. For example, the logic-based plug-in matching of state-based service specifications (PE) can be adopted to the plug-in matching of stateless service signatures (IO): Service S is expected to return more specific output data whose logically defined semantics is equivalent or subsumed by those of the desired output in request R , and requires more generic input data than requested in R .

More concrete, the signature of S plugs into the signature of request R iff $\forall IN_S \exists IN_R: IN_S < IN_R \wedge \forall OUT_R \exists OUT_S: OUT_S \in LSC(OUT_R)$, with $LSC(C)$ the set of least specific concepts (direct children) C' of C , i.e. C' is a immediate sub-concept of C in the shared (matchmaker) ontology. The quantified constraint that S may require less input than specified in R guarantees at a minimum that S is, in principle, executable with the input provided by the user in R . This holds if and only if the logical service input concepts are appropriately mapped to the corresponding WSDL service input message data types in XMLS.

Examples of Semantic Web service matchmakers that perform logic-based semantic matching of service signatures only are the OWLSM [45] and the OWLS-UDDI [81]. Though the latter determines a signature plug-in matching degree which is defined inverse to the original definition and restricted to the output. Approaches to logic-based semantic IOPE matching of Web services are proposed in [48, 103]. In general, logic-based matching of stateless service descriptions with I/O

¹²<http://alice.unibo.it/xwiki/bin/view/Tuprolog/>

concepts and conjunctive constraints on their relationship specified in OWL-DL, that is SHOIN has been proven decidable though intractable [42]. This indicates the respective decidability of service IOPE-matching for OWL-S (with OWL-DL) and WSML (with WSML-DL).

4.2.3 Non-logic-based Semantic Profile Matching

As mentioned above, non-logic-based Semantic Web service matchmaker do not perform any logical inferencing on service semantics. Instead, they compute the degree of semantic matching of given pairs of service descriptions based on, for example, syntactic similarity measurement, structured graph matching, or numeric concept distance computations over given ontologies. There is a wide range of means of text similarity metrics from information retrieval, approximated pattern discovery, and data clustering from data mining, or ranked keyword, and structured XML search with XQuery, XIRQL or TeXQuery [39, 6]. In this sense, non-logic-based semantic service matching means exploit semantics that are implicit in, for example, patterns, subgraphs, or relative frequencies of terms used in the service descriptions, rather than declarative IOPE semantics explicitly specified in the considered logic.

One example is the matchmaker iMatcher1 [14] which imprecisely queries a set of OWL-S service profiles that are stored as serialized RDF graphs in a RDF database with an extension of RDQL, called iRDQL, based on four (token and edit based) syntactic similarity metrics from information retrieval. The imprecise querying of RDF resources with similarity joins bases on TFIDF and the Levenshtein metric. The results are ranked according to the numerical scores of these syntactic similarity measurements, and a user-defined threshold.

The DSD matchmaker [53, 61] performs, in essence, graph matching over pairs of state-based service descriptions in the object oriented service description language DSD (with variables and declarative object sets) without any logic-based semantics. The matching process determines what assignment of IOPE variables is necessary such that the state-based service offer is included in the set (of service instances) defined by the request, and returns a numeric (fuzzy) degree of DSD service matching.

4.2.4 Hybrid Semantic Profile Matching

Syntactic matching techniques are first class candidates for the development of hybrid semantic service profile matching solutions that combine means of both crisp logic-based and non-logic-based semantic matching where each alone would fail. Indeed, first experimental evaluation of the performance of hybrid semantic service matchmakers OWLS-MX and iMatcher2 show that logic-based semantic service selection can be significantly outperformed by the former under certain conditions.

LARKS [105, 105] has been the first hybrid semantic service IOPE matchmaker for services written in a frame-based language called LARKS. The matchmaker OWLS-MX [59] bases in part on LARKS, and is the first hybrid semantic service signature (IO) matchmaker for OWL-S services. OWLS-MX complements deductive (DL) reasoning with approximated IR-based matching. For this purpose, each of its four hybrid variants OWLS-M1 to OWLS-M4 applies a selected token-based string similarity metric (cosine/TFIDF, extended Jaccard, Jensen-Shannon, LOI) to the given pair of service signature strings in order to determine their degree of text similarity-based matching. If the text similarity value exceeds a given threshold the failure of logic-based matching is tolerated, that means the service is eventually classified as semantically relevant to the given query. The ranking aggregates both types of matching degrees with respect to the total order of logic-based matching degrees. Experimental evaluation results over the test collection OWLS-TC together with a FP/FN-analysis of OWLS-MX showed that the performance of logic-based semantic matching can be improved by its combination with non-logic-based text similarity measurement [54, 55].

Similarly, the hybrid semantic service profile matchmaker iMatcher2 [51] uses multiple edit- or token-based text similarity metrics (Bi-Gram, Levenshtein, Monge-Elkan and Jaro similarity measures) to determine the degree of semantic matching between a given pair of OWL-S service profiles. Like OWLS-MX, the iMatcher2 transforms each structured service profile description into a weighted keyword vector that includes not only the names but terms derived by means of logic-based unfolding of its service input and output concepts. In this sense, iMatcher2 classifies as a hybrid matchmaker. The experimental evaluation of iMatcher2 over the test collection OWLS-TC2.1 confirmed, in principle, the previously reported results of the evaluation of OWLS-MX.

In its adaptive mode, iMatcher2 can also be trained over a given retrieval training collection to predict the degree of semantic matching of unknown services to queries by means of selected regression models (support vector regression with a RBF kernel, linear and logistic regression). This regression-based induction is performed over the set of (a) the binary value of subjective semantic relevance as defined in the relevance sets, and (b) different text similarity values computed by means of the selected similarity metrics for each pair of query and service of the training collection. After training, the iMatcher2 first computes the text similarity values (using the selected similarity metrics) of a given query to all services of a given test collection, then uses the learned regression model to predict the combined similarity (or likelihood) of a match, and finally returns the answers in decreasing order of similarity. Experimental evaluation of the adaptive iMatcher2 showed that the combined logical deduction and regression-based learning of text similarities produces superior performance over logical inference only.

The hybrid semantic service matchmaker FC-MATCH [16] does a combined logic-based and text similarity-based matching of single service and query concepts written in OWL-DL. A service concept S is defined as logical conjunction of existential qualified role expressions where each role corresponds to a selected

profile parameter: $S = \exists\text{hasCategory}(C_1) \sqcap \exists\text{hasOperation}(C_2) \sqcap \exists\text{hasInput}(C_3) \sqcap \exists\text{hasOutput}(C_4)$). Hybrid matching degrees are computed by means of (a) combined checking of logic-based subsumption of profile concepts (C_i) and (b) computing the so-called Dice (name affinity) similarity coefficient between terms occurring in these concepts according to the given terminological relationships of the thesaurus WordNet. FC-MATCH (FC stands for functional comparison) performs structured hybrid semantic matching of functional (I/O) and non-functional profile parameters (hasCategory, hasOperation). That is a combined matching of functional and non-functional parameters of OWL-S service profiles rewritten in special OWL-DL expressions. To the best of our knowledge, FC-MATCH has not been experimentally evaluated yet.

WSMO-MX [49] is the first hybrid semantic matchmaker for services written in a WSML-Rule variant, called WSML-MX. The hybrid service matching scheme of WSMO-MX is a combination of ideas of hybrid semantic matching as performed by OWLS-MX, the object-oriented graph matching of the matchmaker DSD-MM, and the concept of intentional matching of services in [48]. WSMO-MX applies different logic-based and text similarity matching filters to retrieve and rank services that are relevant to a query. The hybrid semantic matching degrees are recursively computed by aggregated valuations of (a) ontology-based type matching (logical concept subsumption), (b) logical (instance-based) constraint matching in F-logic(LP) through approximative query containment, (c) relation name matching, and (d) syntactic similarity measurement as well. The experimental evaluation of WSMO-MX over an initial WSML service retrieval test collection is ongoing work.

However, it is not yet known what kind of hybrid service matching will scale best to the size of the Web in practice. Research in this direction is in perfect line with the just recent call in [33] for a general shift in Semantic Web research towards scalable, approximative rather than strict logic-based reasoning.

4.2.5 Logic-based Semantic Process Matching

Semantic matching of service process models, in general, is very uncommon, and not intended by the designers of current Semantic Web Service description formats. Besides, the semantics of process models in OWL-S or WSML have not been formally defined yet, while neither SAWSDL nor monolithic service descriptions offer any process model. This problem can be partly solved by intuitively rewriting the process model descriptions in an appropriate logic with automated proof system and respective analysis tool support.

For example, in [109], OWL-S service process models are mapped into (intuitively) equivalent logical Promela statements that are then efficiently evaluated by the SPIN model checker¹³. This allows to verify the correctness of a given ser-

¹³A model checker verifies if a given system (service process) model satisfies a desirable prop-

vice process model in terms of consistency and liveness properties of an advertised service like the Delivery process always executes after the Buy process. The result of such service process model checking could be used for process-oriented OWL-S service selection (by identifying properties of service process models to be verified with queries to match); this is a topic of ongoing research.

Alternatively, the matching of process models of OWL-S services that are grounded in WSDL (cf. Chapter 3) can be, in principle, reduced to the matching of corresponding WSDL service orchestrations in BPEL. As mentioned before, the OWL-S process model captures a common subset of workflow features that can be intuitively mapped to BPEL (used to define WSDL service compositions) which offers an all-inclusive superset of such features (e.g. structured process activities in BPEL like Assignment, Fault Handler, Terminate are not available in OWL-S) [9]. Though BPEL has been given no formal semantics either yet, there are a few approaches to fill this gap based on Petri nets [69] and abstract state machines [32] that allow to verify liveness properties of WSDL service orchestrations in BPEL [72]. However, there are no approaches to exploit any of the proposed formal BPEL semantics for semantic matching of OWL-S process models that correspond to BPEL orchestrations of WSDL services.

4.2.6 Non-logic-based and Hybrid Semantic Process Model Matching

There are only a few approaches to non-logic-based Semantic Web service process model matching. One approach to the matching of (business) process dependency graphs based on syntactic similarity measurements is presented in [11]. [12] propose a hybrid matchmaker (IO-RPTM) that recursively compares the DAML-S process model dependency graphs based on given workflow operations and logical match between IO parameter concepts of connected (sub-)service nodes of the process graphs. On the other hand, means of functional service process matching can be exploited to search for a set of relevant subservices of a single composite service.

4.2.7 Semantic Service Discovery Architectures

Existing Semantic Web service discovery architectures and systems in the literature can be broadly categorized as centralized and decentralized by the way they handle service information storage and location in the considered service network [5, 37]. A classification of implemented Semantic Web service discovery systems is given in Figure 4.2.

Centralized service discovery systems rely on one single, possibly replicated, global directory service (repository, registry) maintained by a distinguished so called super-peer or middle agent like matchmaker, broker or mediator agent [58].

erty. If the property does not hold, it returns a counter-example of an execution where the property fails.




Figure 4.2: Categories of Semantic Web Service discovery architectures and systems.

Contrary, decentralized service discovery systems rely on distributing service storage information over several peers in a structured, unstructured or hybrid P2P network.

Semantic service discovery systems can be further classified with respect to the kind of semantic service matching means used by the intelligent agents in the network. For example, the exact keyword-based service location mechanisms of all contemporary P2P systems like JINI, SLP, Gnutella flooding, and DHT (distributed hash table) can be complemented or replaced by sophisticated logic-based semantic matching means to improve the quality of the search result.

As mentioned above, due to its generic functionality, any service matchmaker (cf. Figure 4.1) can be used in arbitrary discovery architectures and systems. In the extremes, a matchmaker can either serve as a central service directory (index) or look-up service, or can be integrated into each peer of an unstructured P2P service network to support an informed adaptive service search like in RS2D [13]. In fact, a few means of semantic service matching were originally developed for decentralized semantic P2P service retrieval in different applications.

Centralized Semantic P2P Service Discovery

In centralized semantic P2P service systems, a dedicated central service directory or matchmaker returns a list of providers of semantically relevant services to the requester. Contrary to centralized client-server middleware or brokering, the requester then directly interacts with selected providers for service provision [58]. The advantage of such centralized discovery architectures is a fast resource or service lookup time, though the central look-up server or registry like in JINI or the CORBA-ORB interface registry is a single point of failure that can be only partially mitigated by replication and caching strategies.

An application of centralized P2P service discovery is the Napster music file sharing system, and the SETI@home system that is exploiting a vast set of distributed computational resources world wide to search for extraterrestrial signals. From the Semantic Web Service discovery perspective, each of the above mentioned stand-alone Semantic Web Service matchmakers, in principle, realizes a centralized logic-based semantic service discovery system by itself. For example, the SCALLOPS e-health service coordination system uses the hybrid semantic matchmaker OWLS-MX as a central matchmaker for the selection of relevant e-health services in a medical emergency assistance application. The same matchmaker is distributed to each peer of an unstructured P2P network for decentralized OWL-S service discovery [13].

MWSDI [111] is a centralized semantic P2P service system with non-logic-based semantic service signature matching. Each peer in the system maintains one domain specific WSDL-S (SAWSDL) service registry and respective ontologies; multiple peers can form a domain oriented group. However, a distinguished central gateway or super-peer provides a global registries ontology (GRO) that maintains the complete taxonomy of all domain registries, the mappings between WSDL-S service I/O message types and concepts from shared domain ontologies in the system, associates registries to them, and serves as central look-up service for all peers. This central super-peer is replicated in form of so called auxiliary peers for reasons of scalability. For service location, any client peer (user) selects the relevant domain registries via the central GRO at the super-peer which then performs non-logic-based semantic matching (structural XMLS graph matching, NGram-based syntactic similarity, synonyms/hyponyms/hypernyms in the GRO) of service input and output concepts with those of the desired service. However, it would be hard to build the GRO, and difficult for the user to query the GRO without knowing its details in advance.

Decentralized Semantic P2P Service Discovery

Decentralized semantic service discovery systems rely on service information storage and location mechanisms that are distributed over all peers in structured, unstructured or hybrid P2P networks.

Structured Semantic P2P Service Systems Structured P2P systems have no central directory server but a significant amount of structure of the network topology (overlay) which is tightly controlled. Resources are placed neither at random peers nor in one central directory but at specified locations for efficient querying. In other words, the service index of the system is distributed to all peers according to a given structured overlay enforcing a deterministic content distribution which can be used for routing point queries.

Prominent examples of structured P2P systems are those with flat DHT-based resource distribution and location mechanism like Chord rings, Pastry, Tapestry, CAN, P-Grid and P2PALvis, and structured hierarchic P2P systems. Flat DHT-based systems allow to route queries with certain keys to particular peers containing the desired data. But to provide this functionality all new content in the network has to be published at the peer responsible for the respective key, if new data on a peer arrives, or a new peer joins the network.

In structured hierarchical or N-super-peer P2P systems ($N > 1$), peers are organized in N domain oriented groups with possibly heterogeneous service location mechanisms (e.g hierarchic DHT, that is, one group with Chord ring overlay, another one with P-Grid overlay, etc.). Each group is represented by one super-peer hosting the group/domain service index. The set of super-peers, in turn, can be hierarchically structured with federated service directories in a super-peer top level overlay of the network. Peers within a group query its super-peer which interacts with other super-peers to route the query to relevant peer groups for response. The functionality of a super-peer of one peer group is not necessarily fixed, but, in case of node failure, transferable to a new peer of that group. Typically JXTA, a collection of P2P protocols, is used to realize super-peer based P2P systems, though it does not enforce such architectures.

Examples of decentralized Semantic Web service discovery in structured P2P networks are WSPDS [47], SSLinkNet [66], CASCOS-P2P_{3b} [19], Grid-Vine [1], WSML-P2P [112] and Agora-P2P [60, 67]. SSLinkNet, Agora-P2P and WSML-P2P exploit keyword-based discovery in a Chord ring, respectively, P-Grid system with non-logic-based semantic profile matching of services in WSDL, respectively, WSML. The Grid-Vine system performs non-logic-based semantic P2P content retrieval by means of so-called semantic gossiping with the underlying P-Grid system. The CASCOS and Agora-P2P systems have been demonstrated for logic-based semantic OWL-S (DAML-S) service discovery in hierarchic structured P2P networks.

In the SSLinkNet [66], a Chord ring-based search is complemented by forwarding the same Web service request by the identified peers to relevant neighbors based on a given so-called semantic service link network. The semantic links between services are determined by non-logic-based semantic service matching, and are used to derive semantic relationships between service provider peers based on heuristic rules.

Similarly, the AGORA-P2P system [60, 67] uses a Chord ring as the underlying infrastructure for a distributed storage of information about OWL-S services

over peers. Service input and output concept names are hashed as mere literals to unique integer keys such that peers holding the same key are offering services with equal literals in a circular key space. A service request is characterized as a syntactic multi-key query against this Chord ring. Both systems, SLinkNet and AGORA-P2P, do not cope with the known problem of efficiently preserving the stability of Chord rings in dynamic environments.

The generic CASCOM semantic service coordination architecture has been instantiated in terms of a hierarchic structured P2P network with N interacting super-peers each hosting a domain service registry that make up a federated Web Service directory. Each peer within a group can complement a keyword-based pre-selection of OWL-S services in their super-peer domain registries with a more complex semantic matching by a selected hybrid or logic-based semantic OWL-S matchmaker (ROWL-S, PCEM or OWLS-MX) on demand. Both, the simple service discovery agent and Semantic Web Service matchmaking module are integrated into each peer (cf. Chapter 10).

The Grid-Vine system [1] performs a hybrid semantic search of semantically annotated resources by means of so-called semantic gossiping between peers about their actual semantic knowledge (also called logical layer or semantic overlay of the P2P system). The semantic overlay is defined by (a) the set of peer ontologies in RDFS or XMLS that are used to encode document annotations in RDF (each RDF triple or concept in the peer schema represents a set of documents as its instances), and (b) a set of user-specified peer schema (concept) mappings that are used by the peers to translate received queries. The numeric "semantic quality" value of these directed concept mappings, hence the non-logic-based degree of semantic similarity between query and resource annotation concept of two peers is locally assessed by the requester through a quantitative analysis of (transitive) propagation cycles of the mappings (and their previous semantic quality value) which might be wrong but not by means of logic-based reasoning about concepts. The translation links, that is the mapping and its numeric "semantic quality" are continuously exchanged and updated by the peers: Semantic gossiping among peers is the propagation of queries to peers for which no direct but transitive translation links exist. The efficient location of resources for a given and translated query by the underlying P-Grid system bases on keyword-based matching of their identifiers, that are DHT keys.

Service discovery in structured P2P networks can provide search guarantees, in the sense of total service recall in the network, while simultaneously minimizing messaging overhead. Typically, structured networks such as DHT-based P2P networks of n peers offer efficient $O(\log(n))$ search complexity for locating even rare items, but they incur significantly higher organizational overheads (maintaining DHT, publishing)¹⁴ than unstructured P2P networks. Alternatively, flooding-based or random-walks discovery in unstructured P2P networks are effective for

¹⁴For example, peer p_n publishes each of its hashed items ($term_i$) over the DHT network, that is the item gets stored in an inverted list ($term_i, [..., p_n, ...]$) of some peer that is found in $O(\log(n))$ hops.

locating highly replicated, means popular, but not rare items. Hybrid designs of P2P networks aim to combine the best of both worlds such as using random-walks (with state-keeping to prevent walkers from revisiting same peers) for locating popular items, and structured (DHT) search techniques for locating rare items [70].

Unstructured Semantic P2P Service Systems In unstructured P2P systems, peers initially have no index nor any precise control over the network topology (overlay) or file placement based on any knowledge of the topology. That is, they do not rely on any structured network overlay for query routing as they have no inherent restrictions on the type of service discovery they can perform.

For example, resources in unstructured P2P systems like Gnutella or Morpheus are located by means of network flooding: Each peer broadcasts a given query to all neighbour peers within a certain radius (TTL) until a service is found, or the given query TTL is zero. Such network flooding is extremely resilient to network dynamics (peers entering and leaving the system), but generates high network traffic.

This problem can be mitigated by a Random Walk search where each peer builds a local index about available services of its direct neighbour peers over time and randomly forwards a query to one of them in DFS manner until the service is found¹⁵ as well as replication and caching strategies based on, for example, access frequencies and popularity of services [71].

Approaches to informed probabilistic adaptive P2P search like in APS [108] improve on such random walks based on estimations over dynamically observed service location information stored in the local indices of peers. In contrast to the structured P2P search, this only provides probabilistic search guarantees, that is incomplete recall.

In any case, the majority of unstructured P2P service systems only performs keyword-based service matching and does not exploit any qualitative results from logic-based or hybrid semantic service matching to improve the quality of an informed search. In fact, only a few system are available for logic-based or hybrid Semantic Web service retrieval such as DReggie/GSD [22, 23], HyperCuP [95], Sem-WSPDS [47], [82], Bibster [40], INGA [68], and RS2D [13]. These systems differ in the way of how peers perform flooding or adaptive query routing based on evolving local knowledge about the semantic overlay, that is knowledge about the semantic relationships between distributed services and ontologies in unstructured P2P networks. Besides, all existing system implementations, except INGA and Bibster, perform semantic service IO profile matching for OWL-S (DAML-S), while HyperCuP peers dynamically build a semantic overlay based on monolithic service concepts.

¹⁵This is valid in case the length of the random walk is equal to the number of peers flooded with bounded TTL or hops).

For example, [82] proposes the discovery of relevant DAML-S services in unstructured P2P networks based on both the Gnutella P2P discovery process and a complementary logic-based service matching process (OWLS-UDDI matchmaker) over the returned answer set. However, the broadcast or flooding-based search in unstructured P2P networks like Gnutella is known to suffer from traffic and load balancing problems.

Though Bibster and INGA have not been explicitly designed for Semantic Web service discovery, they could be used for this purpose. In INGA [68], peers dynamically adapt the network topology, driven by the dynamically observed history of successful or semantically similar queries, and a dynamic shortcut selection strategy, which forwards queries to a community of peers that are likely to best answer given queries. The observed results are used by each peer for maintaining a bounded local (recommender) index storing semantically labelled topic specific routing shortcuts (that connect peers sharing similar interests).

Similarly, in Bibster [40] peers have prior knowledge about a fixed semantic overlay network that is initially built by means of a special first round advertisement and local caching policy. Each peer only stores those advertisements that are semantically close to at least one of their own services, and then selects for given queries only those two neighbours with top ranked expertise according to the semantic overlay it knows in prior. Further, prior knowledge about other peers ontologies as well as their mapping to local ontologies is assumed. This is similar to the ontology-based service query routing in HyperCuP [95].

In RS2D [13], contrary to Bibster and DReggie/GSD, the peers perform an adaptive probabilistic risk-driven search for relevant OWL-S services without any fixed or prior knowledge about the semantic overlay. Each peer uses an integrated OWLS-MX matchmaker for hybrid semantic IO matching of local services with given query, and dynamically learns the average query-answer behaviour of its direct neighbours in the network. The decision to whom to forward a given semantic service request is then driven by the estimated mixed individual Bayes' conditional risk of routing failure in terms of both semantic loss and high communication costs. Peers are dynamically maintaining their local service (matchmaker) ontology-based on observations of the results which, in particular, renders RS2D independent from the use of any fixed global ontology for semantic annotation like in DReggie/GSD.

Semantic Hybrid P2P Service Systems Hybrid P2P search infrastructures combine both structured and unstructured location mechanisms. For example, Edutella combines a super-peer network with routing indices and an efficient broadcast. In [70] a flat DHT approach is used to locate rare items, and flooding techniques are used for searching highly replicated items. A similar approach of hybrid P2P query routing that adaptively switches between different kinds of structured and unstructured search together with preliminary experimental results are reported in [94]. However, there are no hybrid P2P systems for semantic service discovery

available yet.

Despite recent advances in the converging technologies of semantic Web and P2P computing [102], the scalability of semantic service discovery in structured, unstructured or hybrid P2P networks such as those for real-time mobile ad-hoc network applications is one major open problem. Research in this direction has just started. Preliminary solutions to this challenge vary in the expressivity of semantic service description, and the complexity of semantic matching means ranging from computationally heavy Semantic Web Service matchmakers like OWLS-MX in SCALLOPS and CASCOM, to those with a streamlined DL reasoner such as Krhype [52] suitable for thin clients on mobile devices in IASON [35]. An example analysis of semantic service discovery architectures for realizing a mobile e-health application is given in [21].

4.3 Semantic Service Composition Planning

Semantic Web service composition is the act of taking several semantically annotated component services, and bundling them together to meet the needs of a given customer. Automating this process is desirable to improve speed and efficiency of customer response, and, in the semantic Web, supported by the formal grounding of service and data annotations in logics.

4.3.1 Web Service Composition

In general, Web service composition is similar to the composition of workflows such that existing techniques for workflow pattern generation, composition, and management can be partially reused for this purpose [41]. Typically, the user has to specify an abstract workflow of the required composite Web service including both the set of nodes (desired services) and the control and data flow between these nodes of the workflow network. The concrete services instantiating these nodes are bound at runtime according to the abstract node descriptions, also called “search recipes” [20]. In particular, the mainstream approach to composition is to have a single entity responsible for manually scripting such workflows (orchestration and choreography) between WSDL services of different business partners in BPEL [83, 2]. This is largely motivated by industry to work for service composition in legally contracted business partner coalitions — in which there is, unlike in open service environment, only very limited need for automated service composition planning, if at all. Besides, neither WSDL nor BPEL or any other workflow languages like UML2 or YAWL have formal semantics which would allow for an automated logic-based composition.

In fact, the majority of existing composition planners for Semantic Web services draws its inspiration from the vast literature on logic-based AI planning [84]. In the following, we focus on these approaches to Semantic Web service composition, and comment on the interleaving of service composition planning with

discovery, and distributed plan execution. Please note that, the set of presented examples of Semantic Web Service composition planners is representative but not exhaustive.

4.3.2 AI-Planning-Based Web Service Composition

The service composition problem roughly corresponds to the state-based planning problem (I, A, G) in AI to devise a sound, complete, and executable plan which satisfies a given goal state G by executing a sequence of services as actions in A from a given initial world state I . Classical AI planning-based (planning from first principles) composition focuses on the description of services as merely deterministic state transitions (actions) with preconditions and state altering (physical) effects. Actions are applicable to actual world states based on the evaluation of preconditions and yield new (simulated) states where the effects are valid. Further, classical AI planning is performed under the assumption of a closed world with complete, fully observable initial (world) state such that no conditional or contingency planning under uncertainty is required. This is not necessarily appropriate for service composition planning in the dynamic and open-ended Semantic Web [101] where (a) the initial state can be incomplete, and actions may have several possible (conditional) outcomes and effects modeled in the domain but not known at design time. However, all existing SWS composition planners are closed-world planners of which some are able to cope with uncertainties about the domain.

A given logical goal expression and set of logic-based definitions of semantic service signature (I/O) concepts together with logic preconditions and effects from a DL-based ontology (domain or background theory) can be converted into one declarative (FOL) description of the planning domain and problem - which can serve a given logic-based AI planner as input. In particular, service outputs are encoded as special non-state altering knowledge effects, and inputs as special preconditions. The standard target language for the conversion is PDDL (Planning Domain Description Language) but alternative representation formalisms are, for example, the situation calculus, linear logic [92], high-level logic programming languages based on this calculus like GOLOG [73], Petri nets, or HTN planning task and method description format [99].

In the following, we classify existing Semantic Web service composition planners and comment on the principled interrelation between composition, discovery, and execution. Please note that the set of presented examples is representative but not exhaustive.

4.3.3 Classification of Semantic Service Composition Planners

In general, any AI planning framework for Semantic Web service composition can be characterized by

- the representation of the planning domain and problem to allow for automated reasoning on actions and states,
- the planning method applied to solve the given composition problem in the domain, and
- the parts of service semantics that are used for this purpose.

We can classify existing Semantic Web service composition planners according to the latter two criteria, which yields the following classes.

- Dynamic or static Semantic Web service composition planners depending on whether the plan generation and execution are inherently interleaved in the sense that actions (services) can be executed at planning time, or not.
- Functional-level or process-level Semantic Web service composition planners depending on whether the plan generation relies on the service profile semantics only, or the process model semantics in addition (data and control flow) [63].

Figure 4.3 shows the respective classification of existing Semantic Web service composition planners.

Static and Dynamic Composition

The majority of Semantic Web service composition planners such as GOAL [86], MetaComp (cf. Chapter 11), PLCP [89], RPCLM-SCP [63] and AGORA-SCP [92] are static classical planners. Approaches to dynamic composition planning with different degrees of interleaving plan generation and execution are rare. Unlike the static case, restricted dynamic composition planners allow the execution of information gathering but no world state altering services, hence are capable of planning under uncertainty about action outcomes at planning time. Examples of such composition planners are SHOP2 [97, 99], GOLOG-SCP [73] and OWLS-XPlan1 [56].

Advanced and reactive dynamic composition planners in stochastic domains even take non-deterministic world state changes into account during planning. While advanced dynamic planners like OWLS-XPlan2 [57] are capable of heuristic replanning subject to partially observed (but not caused) state changes that affect the current plan at planning time, their reactive counterparts like INFRAWEB-RTC [4] fully interleave their plan generation and execution in the fashion of dynamic contingency and real-time planning.

Functional- and Process-Level Composition

As shown in Figure 4.3, most Semantic Web Service composition planners perform functional-level or service profile-based composition (FLC) planning. FLC planning considers services as atomic or composite black-box actions which functionality can solely be described in terms of their inputs, outputs, preconditions,

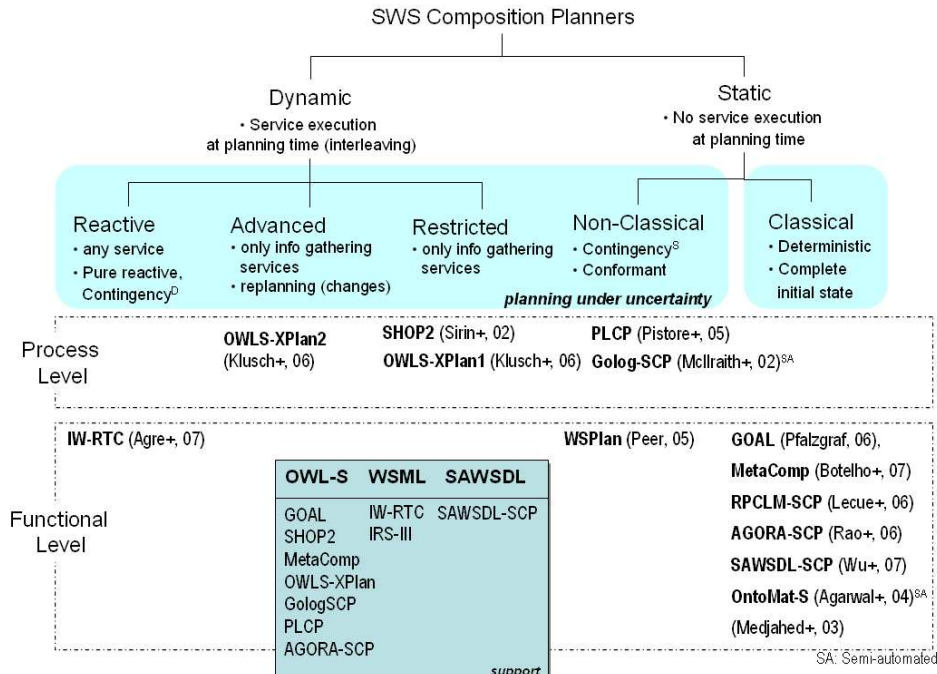


Figure 4.3: Classes of Semantic Web Service composition planners.

and effects, and which can be executed in a simple request-response without interaction patterns. Examples of FLC planners are GOAL [86], SAWSDL-SCP [113] and OntoMat-S [3].

Process-level composition (PLC) planning extends FLC planning in the sense that it also the internal complex behavior of existing services into account. Prominent examples are SHOP2 [99], PLCP [87, 89] and OWLS-XPlan [56, 57]. Both kinds of composition planning perform, in particular, semantic service profile or process matching which is either inherent to the underlying planning mechanism, or achieved by a connected stand-alone Semantic Web service matchmaker. We will discuss the interrelation between composition and semantic matching later.

Support of Semantic Web Service Description Frameworks

Remarkably, most implemented Semantic Web Service composition planners support OWL-S like GOAL, OWLS-XPlan, SHOP2, GologSCP and MetaComp, while there is considerably less support of the standard SAWSDL and WSML available to date. In fact, the SAWSDL-SCP planner [113] is the only one for SAWSDL, while the IW-RTC planner [4] is, apart from the semi-automated orchestration of

WSML services in IRS-III, the only fully automated FLC planner for WSML yet.

Most composition planner feature an integrated conversion of Semantic Web Services, goals and ontologies into the internally used format of the planning domain and problem description, though a few others like the framework WSPan [85] for static PDDL-based planning under uncertainty, and the recursive, progression-based causal-link matrix composition planner RPCLM-SCP [63] do not.

In the following, we discuss each category and selected examples of Semantic Web service composition planners in more detail.

4.3.4 Functional-Level Composition Planners

Intuitively, FLC planning generates a sequence of Semantic Web Services based on their profiles that exact or plug-in matches with the desired (goal) service. In particular, existing services S_i, S_{i+1} are chained in this plan such that the output of S_i matches with the input of S_{i+1} , while the preconditions of S_{i+1} are satisfied in the world state after execution of S_i . Depending on the considered Semantic Web Service description format (cf. Chapter 3), different approaches to logic-based, non-logic-based or hybrid semantic service profile IOPE matching are available for this purpose (cf. Figure 4.1).

In order to automatically search for a solution to the composition problem, FLC planners can exploit different AI planning techniques with inherent logic-based semantic profile IOPE- or PE-matching like WSPan [85], respectively, MetaComp (cf. 11). The recursive forward-search planner GOAL [86] as well as the SAWSDL-SCP [113] apply non-logic-based semantic profile IO matching of OWL-S, respectively, SAWSDL services.

In AGORA-SCP [92], theorem proving with hybrid semantic profile IO matching is performed for OWL-S service composition: Both services and a request (theorem) are described in linear logic, related to classical FOL, while the SNARK theorem prover is used to prove that the request can be deduced from the set of services. The service composition plan then is extracted from the constructive proof.

The FLC planner in [75] uses proprietary composability rules for generating all possible plans of hybrid semantic profile IO matching services in a specific description format (CSSL). From these plans the requester has to select the one of best quality (QoS).

4.3.5 Process-Level Semantic Service Composition Planners

Though FLC planning methods can address conditional outputs and effects of composite services with dynamic planning under uncertainty, considering services as black-boxes does not allow them to take the internal complex service behaviour into account at planning time. Such behavior is usually described as subservice interactions by means of control constructs including conditional and iterative

steps. This is the domain of process level composition (PLC) planning that extends FLC planning in the aforementioned sense.

However, only few approaches to process level composition planning for Semantic Web Services exist to date. For example, orchestration of WSML services in IRS-III [31] synthesizes abstract state machines to compose individual services in a given process flow defined in OCML¹⁶. Though, the functionality of the WSMX orchestration unit has not been completely defined yet.

Other automated PLC planners of OWL-S services exploit different AI planning techniques such as

- HTN (Hierarchical Task Network) planning of OWL-S process models converted to HTN methods like in SHOP2 [99],
- Neo-classical GRAPHPLAN-based planning mixed with HTN planning of OWL-S services converted to PDDL in OWLS-XPlan [56, 57],
- Value-based synthesis of OWL-S process models in a given plan template of situation calculus-based GOLOG programs [73, 74],
- Planning as model checking of OWL-S process models converted to equivalent state transition systems (STS) in the PLCP [87, 89].

In the following, we discuss each class of static and dynamic Semantic Web service composition planners together with selected examples, if available.

4.3.6 Static Semantic Service Composition Planners

The class of static AI planning-based composition covers approaches to both classical and non-classical planning under uncertainty.

Static Classical Planning

As mentioned above, classical AI planners perform (off-line) planning under the assumption of a closed, perfect world with deterministic actions and a complete initial state of a fully observable domain at design time. For example, Graphplan is a prominent classical AI planning algorithm that first performs a reachability analysis by constructing a plan graph, and then performs logic-based goal regression within this graph to find a plan that satisfies the goal. Classical AI planners are static since their plan generation and execution is strictly decoupled.

Examples of Static Classical Composition Planners

One example of a static classical Semantic Web service composition planner is GOAL [86] developed in the SmartWeb project. GOAL composes extended OWL-S services by means of a classical recursive forward-search [36]. Both, the initial

¹⁶kmi.open.ac.uk/projects/ocml/

state and the goal state are derived from the semantic representation of the user's question (goal) obtained by a multimodal dialogue system in SmartWeb. At each stage of the planning process the set of services which input parameters are applicable to the current state is determined by signature (IO) matching through polynomial subgraph isomorphism checking [76]: The instance patterns of input parameters are matched against the graph representation of the state, and a service is applied to a plan state (simulated world state) by merging the instance patterns of its output parameters with the state. As a result, GOAL does not exploit any logical concept reasoning but structural service I/O graph matching to compose services. If plan generation fails, GOAL detects non-matching paths within instance patterns and consequently produces a clarification request (aka information gathering service) conveyed to the user by the dialogue system; on response by the user the planning process is restarted in total.

Static service composition in the AGORA-SCP service composition system [92] relies on linear logic (LL) theorem proving. The profiles of available DAML-S services are translated in to a set of LL axioms, and the service request is formulated as a LL theorem to be proven over this set. In case of success, the composition plan can be extracted from the proof, transformed to a DAML-S process model and executed as a BPEL script. The AGORA planner is the only approach to decentralized composition planning in structured P2P networks [60].

An example of a static classical Semantic Web Service composition planner based on a special logic-based PDDL planner is MetaComp which we describe in detail in Chapter 11.

Static Planning under Uncertainty

Work on planning under uncertainty in AI is usually classified according to (a) the representation of uncertainty, that is whether uncertainty is modeled strictly logically, using disjunctions, or is modeled numerically (e.g. with probabilities), and (b) observability assumptions, that is whether the uncertain outcomes of actions are not observable via sensing actions (conformant planning); partially or fully observable via sensing actions (conditional or contingency planning) [26]. As mentioned above, we can have uncertainty in the initial states and in the outcome of action execution. Since the observation associated to a given state is not unique, it is also possible to model noisy sensing and lack of information. Information on action outcomes or state changes that affect the plan can be gathered either at planning time (dynamic) or thereafter (static) for replanning purposes.

Static Conditional or Contingency Planning. Static conditional or contingency planner like Cassandra and DTPOP devise a plan that accounts for each possible contingency that may arise in the planning domain. This corresponds to an optimal Markov policy in the POMDP framework for planning under uncertainty

with probabilities, costs and rewards over a finite horizon. The contingency planner anticipates unexpected or uncertain outcomes of actions and events by means of planned sensing actions, and attempts to establish the goals for each different outcome of these actions through conditional branching of the plan in advance¹⁷. The plan execution is driven by the outcome of the integrated sensing subplans for conditional plan branches, and decoupled from its generation which classifies these planners as static.

Static Conformant planning. Conformant planners like the Conformant-FF, Buridan, and UDTPOP perform contingency planning without sensing actions. The problem of conformant planning to search for the best unconditional sequence of actions under uncertainty of initial state and action outcome can be formalized as fully non-observable MDP, as a particular case of POMDP, with a search space pruned by ignoring state observations in contingency planning. For example, conformant Graphplan planning (CGP) [100] expresses the uncertainty in the initial state as a set of completely specified possible worlds, and generates a plan graph for each of these possible worlds in parallel. For actions with uncertain outcomes the number of possible worlds is multiplied by the number of possible outcomes of the action. It then performs a regression (backward) search on them for a plan that satisfies the goal in all possible worlds which ensures that the plan can be executed without any sensory actions. Conformant planners are static in the sense that no action is executed at planning time.

Examples of Static Composition Planners under Uncertainty

The PLCP [88, 89] performs static PLC planning under uncertainty for OWL-S services. OWL-S service signatures and process models together with a given goal are converted to non-deterministic and partially observable state transition systems which are composed by a model checking-based planner (MBP)[87] to a new STS which implements the desired composed service. This STS eventually gets transformed to an executable service composition plan (in BPEL) with possible conditional and iterative behaviors. No action is executed at planning time, and uncertainty is resolved by sensing actions during plan execution.

An example of static FLC planning under uncertainty is the WSPlan framework [85] which provides the user with the option to plug in his own PDDL-based planner and to statically interleave planning (under uncertainty) with plan execution. Static interleaving refers to the cycle of plan generation, plan execution, and replanning based on the result of the executed sensing subplans (in the fashion

¹⁷Examples of decision criteria according to which contingency branches are inserted in the (conventional) plan, and what the branch conditions should be at these points, are the maximum probability of failure, and the maximum expected future reward (utility) as a function of, for example, time and resource consumption. Uncertainty is often characterized by probability distributions over the possible values of planning domain predicates.

of static conditional planning) until a sequential plan without sensing actions is generated that satisfies the goal. There are no static classical PLC planner for Semantic Web Services with deterministic (sequential) process models of composite services only available.

4.3.7 Dynamic Composition Planners

The class of dynamic AI-planning-based composition covers approaches to restricted, advanced and reactive dynamic planning under uncertainty.

Restricted Dynamic Planning

Dynamic planning methods allow agents to inherently interleave plan generation and execution. In restricted dynamic planning, action execution at planning time is restricted to information gathering (book-keeping callbacks) about uncertain action outcomes. These special actions add new knowledge in form of ground facts to the partial observable initial state under the known IRP (Invocation and Reasonable Persistence) assumption [73] to ensure conflict avoidance¹⁸. Like in classical planning, however, world state altering services with physical effects (in opposite to knowledge effects of service outputs) are only simulated in local planning states and never get executed at planning time.

Examples of Restricted Dynamic Composition Planners

Prominent examples of restricted dynamic composition planners are SHOP2, and OWLS-XPlan1 [56] for OWL-S services of which we describe the latter in detail in Chapter 11. SHOP2 [97, 98] converts given OWL-S service process models into HTN methods and applies HTN-planning interleaved with execution of information gathering actions to compose a sequence of services that satisfies the given task. By mapping any OWL-S process model to a situation calculus-based GOLOG program, the authors prove that the plans produced are correct in the sense that they are equivalent to the action sequences found in situation calculus. HTN planning is correct and complete but undecidable due to possibly infinite recursive decomposition of given methods to executable atomic tasks. SHOP2 detects and breaks such decomposition cycles.

Advanced Dynamic Planning

Advanced dynamic planning methods allow in addition to react on arbitrary changes in the world state that may affect the current plan already during planning

¹⁸The IRP assumption states that (a) the information gathered by invoking the service once cannot be changed by external or subsequent actions, and (b) remains the same for repeating the same call during planning. That is, the incremental execution of callbacks would have the same effect when executing them prior to planning in order to complete the initial state for closed world planning.

such as in OWLS-XPlan2. This is in contrast to static planning under uncertainty where sensing subplans of a plan are executed at run time only. However, in both restricted and advanced dynamic planning the interleaved execution of planning with world state altering services is prohibited to prevent obvious problems of planning inconsistencies and conflicts.

Examples of Advanced Dynamic Composition Planners

To the best of our knowledge, OWLS-XPlan2 [57] still is the only one implemented example of an advanced dynamic composition planner. OWLS-XPlan2 will be described in Chapter 11.

Reactive Dynamic Planning

Finally, reactive dynamic planning like in Brooks's subsumption architecture, RETE-based production rule planners, and the symbolic model checking-based planner SyPEM [15] allows the execution of arbitrary actions at planning time. Pure reactive planners produce a set of condition-action (if-then) or reaction rules for every possible situation that may be encountered, whether or not the circumstances that would lead to it can be envisaged or predicted. The inherently interleaved planning and execution is driven through the evaluation of state conditions at every single plan step to select the relevant if-then reaction rule and the immediate execution of the respective, possibly world state altering action; This cycle is repeated until the goal is hopefully reached.

A variant of reactive dynamic planning is dynamic contingency planning like in XII and SAGE. In this case, a plan that is specified up to the information-gathering steps gets executed to that stage, and, once the information has been gathered, the rest of the plan is constructed. Interleaving planning and execution this way has the advantage that it is not necessary to plan for contingencies that do not actually arise. In contrast to pure reactive planners, reasoning is only performed at branch points predicted to be possible or likely.

In any case, reactive dynamic planning comes at the possible cost of plan optimality, and even plan existence, that is suboptimality and dead-end action planning or failure. The related ramification problem¹⁹ is usually addressed either by restrictive assumptions on the nature of service effects on previous planning states [15] in safely explorable domains, or by integrated belief revision (TMS) in the planners knowledge base at severe computational costs.

Examples of Reactive Dynamic Composition Planners

One example of an implemented reactive dynamic composition planner is the real-time composition planner IW-RTC [4] developed in the European research project

¹⁹The problem of ensuring the consistency of the planners knowledge base and the reachability of the original goal in spite of (highly frequent) world state altering service execution during plan generation.

INFRAWEBs. It successively composes pairs of keyword-based IO matching services, executes them and proceeds with planning until the given goal is reached. Unfortunately, the authors do not provide any detailed description of the composition and matching process nor complexity analysis.

Problems of Composition Planning under Uncertainty

One problem with adopting planning under uncertainty for semantic service composition is that the execution of information gathering (book keeping) or even world state altering services at design or planning time might not be charge free, if granted by providers at all. That is, the planning agent might produce significant costs for its users even without any return value in case of plan generation or execution failure. Another problem is the known insufficient scalability of conditional or conformant planning methods to planning domains at Web scale or business application environments with potentially hundreds of thousands of services and vast instance bases. Research on exploiting conditional or conformant planning methods for Semantic Web Service composition has just started.

4.3.8 FLC Planning of Monolithic DL-Based Services

Research on AI-based FLC planning with monolithic DL-based descriptions of services has just started. Intuitively, the corresponding AI planning (plan existence) problem for the composition of such services is as follows. Given an acyclic TBox T describing the domain or background theory in a DL, ABoxes S and G which interpretations I (consistent wrt T) over infinite sets of individual (object) names are describing, respectively the initial and goal state, and a set A of operators describing deterministic, parameterized actions α which precondition and effects are specified in the same DL and transform given interpretations of concepts and roles in T ($I \rightarrow_{\alpha}^T I'$), is there a sequence of actions (consistent with T)²⁰ obtained by instantiating operators with individuals which transforms S into G ?

It has been shown in [10] that the standard reasoning problems on actions, that are executability²¹ and projection²², are decidable for description logics between ALC and ALCOIQ. Furthermore, it has been shown in [77] only recently that the plan existence problem for such actions in ALCOIQ is co-NEXPTIME decidable for finite sets of individuals used to instantiate the actions, while it is known to be PSPACE-complete for propositional STRIPS-style actions. In addition, the extended plan existence problem with infinitely countable set of individuals was proven undecidable, as it is for Datalog STRIPS actions, for actions specified in ALC_U with universal role U for assertions over the whole domain by reduction

²⁰An action is consistent with TBox T , if for every model I of T there exists I' s.t. $I \rightarrow_{\alpha}^T I'$.

²¹Action executability is equal to the satisfaction of action preconditions in given world states: $I \models pre_1, \forall i, 1 \leq i \leq n, I'.I \rightarrow_{\alpha_1 \dots \alpha_i}^T I' : I' \models pre_{i+1}$.

²²Satisfaction of assertion ϕ as a consequence or conjunctive effect of applying actions to a given state: For all models I of S and $T, I'.I \rightarrow_{\alpha_1 \dots \alpha_n}^T I' : I' \models \phi$

to the halting problem of deterministic Turing machines. However, there is no implemented composition planner for monolithic DL-based services available to date.

4.4 Interrelations

In the following, we briefly comment on the principled relations between semantic service composition planning, discovery, and execution. Selected approaches to interleaved semantic service composition planning with negotiation are presented in the introduction to the next part of this thesis.

Semantic Web service composition planning and discovery From the view of semantic service discovery, the composition of complex services is of importance if no available service satisfies the given request. In this case, the matchmaker or requester agent can interact with a composition planner to successfully generate a composite service that eventually satisfies the query.

On the other hand, semantic service composition planning agents require a description of the planning domain and goal to start their planning. Both can be semi-automatically generated from the set of available semantic service descriptions together with related logic-based ontologies, the so-called background theories. In fact, from the view of composition planning, semantic service discovery is of importance for the following reasons: A semantic service matchmaker can be used to

- Prune the initial search space of the composition planner with respect to given application-specific preferences of available services, and
- Select semantically equivalent or plug-in, and execution compatible services during planning as alternative (substitute) services in case of planning failures (replanning).

There is no agreed-upon strategy for pruning the search space of Semantic Web service composition planners. Such pre-filtering of services by a matchmaker can be heuristically performed against non-functional and functional service semantics in order to speed up the corresponding planning process - but at the cost of its incompleteness. That is, composition planning over heuristically pruned search space does not, in general, solve the plan existence problem.

Another source of the same problem, that is correct but incomplete composition planning is the naive interleaving of planning with semantic service matching. For example, the sequential composition of stateful services from a given initial state by consecutive calls of a logic-based semantic service matchmaker by the planner only does not guarantee to find a solution if it exists: Any (not specific planning-oriented) matchmaker usually

- does not maintain any planning state information, thus ignores variable bindings that hold for service signatures (IO) and specifications (PE) according to the actual state reached by the calling (closed-world) planner, and
- performs pairwise service matching only, hence would not return services to the calling planner which combined effects (even with provided state-based instantiation) would eventually lead to a solution.

To the best of our knowledge, all available Semantic Web service matchmakers (cf. introduction to part two) are implemented as a stand-alone tool for mere semantic service matching without any composition planning support. However, functional-level composition planning is a kind of state-based semantic plug-in matching of the generated service plan with the given goal: Any FLC-planner generates a sequence of Semantic Web services based on their profiles that exact or plug-in matches with the desired (goal) service, whereas for each consecutive pair of planned services S and S' the output of S semantically matches with the input of S' , and the preconditions of S' are satisfied in the planning state including the effects of S .

Examples There are only a few implemented approaches that explicitly interleave semantic matching with composition planning.

In [64], logic-based service matching is extended with concept abduction to provide explanations of mismatches between pairs of service profiles that are iteratively used as constructive feedback during composition planning and replanning when searching for alternative services to bridge identified semantic gaps between considered IOPE profiles of services in the current plan step. A similar abduction-based matchmaking approach is presented in [29]. This scenario of explicitly interleaved discovery and composition has been implemented and tested in a non-public France Telecom research project.

In [61], the functional level composition of services specified in the DIANE service description language DSD is explicitly integrated with a DSD matchmaker module that matches service requests asking for multiple connected effects of configurable services. By using a value propagation mechanism and a cut of possible (not actual) parameter value fillings for service descriptions that cover multiple effects the authors avoid exponential complexity for determining an optimal configuration of plug-in matching service advertisements used for a composition.

In [17], the syntactic functional level service composition is based on partial matching of numerically encoded service IO data types in a service directory. Unfortunately, the justification of the proposed numeric codings for matching services appears questionable, though it was shown to efficiently work for certain applications.

The composition planner OWLS-XPlan2 [93] integrates planning-specific service IOPE matching on the grounding level: At each plan step, the planner calls the component OWLS-MXP of the matchmaker OWLS-MX 1.1 to check the compatibility of XMLS types of input and output parameters of consecutive services.

This ensures the principled executability of the generated sequential plan at the service grounding level in WSDL.

The interactive OWL-S service composer developed at UMBC [98] uses the OWLS-UDDI matchmaker to help users filter and select relevant services while building the composition plan. At each plan step, the composer provides the user with advertised services which signatures (IO) plug-in or exact match with that of the last service in the current plan. This leads to an incremental forward chaining of services which does not guarantee completeness without respective user intervention.

The Agora-P2P service composition system [60] is the only approach to decentralized Semantic Web Service composition planning. It uses a Chord ring to publish and locate OWL-S service descriptions keyword-based while linear logic theorem proving and logic-based semantic service IO matching is applied to compose (and therefore search for relevant subservices of) the desired service.

4.4.1 Composition Planning and Execution

The semantic compatibility of subsequent services in a plan does not guarantee their correct execution in concrete terms on the grounding level. A plan is called correct, if it produces a state that satisfies the given goal [63]. The principled plan executability, also called execution composability of a plan requires its data flow to be ensured during plan execution on the service grounding level [75]. This can be verified through complete (XMLS) message data type checking of semantically matching I/O parameters of every pair of subsequent services involved in the plan. For example, OWLS-XPlan2 calls a special matchmaker module that checks plan execution compatibility at each plan step during planning.

The consistent, central or decentral plan execution can be achieved by means of classical (distributed) transaction theory and systems. An advanced and implemented approach to distributed Semantic Web Service composition plan execution is presented, for example, in Chapter 12 (Semantic Web Service Execution) and [79]. However, the availability of non-local services that are not owned by the planning agent can be, in principle, refused by autonomous service providers without any prior commitment at any time. This calls for effective replanning based on alternative semantic matching services delivered by the matchmaker to the composition planner prior to, or during planning such as in OWLS-XPlan2.

4.4.2 Negotiation

Services may not be for free but pay per use. In particular, requester agents might be charged for every single invocation of services at discovery or planning time. Besides, the service pricing is often private which makes it hard, if not infeasible, for any search or composition agent to determine the total expenses of coordinated service value provision to its user.

Standard solution is to negotiate service level agreements and contracting of relevant services based on non-functional service parameters such as QoS, pricing, and reputation between service requester and provider agents involved [114]. Usually, such negotiation takes place after service discovery depending on service configurations and user preferences, followed by contracting [90]. Most existing Semantic Web service frameworks offer slots for non-functional provenance information as part of their service description.

However, the problem of how to dynamically interleave composition (re-)planning and negotiation remains open. Related work draw upon means of parallel auctioning [91], and coalition forming [80] of planning agents in different competitive settings.

4.5 Open Problems

The research field of Semantic Web service coordination is in its infancies. Hence, it comes for no surprise that there are many open problems of both semantic service discovery and composition planning that call for intensive further investigation in the domain. Some major open problems of semantic service discovery are the following.

- *Approximated matching.* How to deal with uncertain, vague or incomplete information about the functionality of available services and user preferences for service discovery? Fuzzy, probability, and possibility theory are first class candidates for the design of approximated (hybrid) semantic service matching algorithms to solve this problem. In particular, efficient reasoners for respective extensions of semantic Web (rule) languages like probabilistic pOWL, fuzzyOWL, or pDatalog can be applied to reason upon semantic service annotations under uncertainty and with preferences.

However, there are no such semantic service matchmakers available yet. Apart from the first hybrid matchmakers for OWL-S and WSML services, OWLS-MX and WSMO-MX, the same holds for the integrated use of means of statistical analysis from data mining or information retrieval for approximative matching of semantic service descriptions.

- *Scalability.* How to reasonably trade off the leveraging of expensive logic-based service discovery means with practical requirements of resource bounded, just-in-time and light-weight service discovery in mobile ad-hoc or unstructured P2P service networks? What kind of approximated and/or adaptive semantic service discovery techniques scale best for what environment (network, user context, services distribution, etc) and application at hand? The required very large scale, comparative performance experiments under practical real-world conditions have not been conducted yet.
- *Adaptive discovery.* How to leverage semantic service discovery by means of machine learning and human-agent interaction? Though a variety of adaptive

personal recommender and user interface agents have been developed in the field, none of the currently implemented semantic Web Service matchmakers is capable of flexibly adapting to its changing user, network, and application environment.

- *Privacy.* How to protect the privacy of individual user profile data that are explicit or implicit in service requests submitted to a central matchmaker, or relevant service providers? Approaches to privacy preserving Semantic Web Service discovery are still very rare, and research in this direction appears somewhat stagnant. Amongst the most powerful solutions proposed are the Rei language for annotating OWL-S services with privacy and authorization policies [28, 46], and the information flow analysis based checking of the privacy preservation of sequential OWL-S service plans [43, 44]. However, nothing is known about the scalability of these solutions in practice yet.
- *Lack of tool support and test collections.* Current easy to use tool support of Semantic Web Service discovery is still lagging behind the theoretical advancements, though there are differences to what extent this is valid for what service description framework (cf. Figure 4.1). In particular, there is no official test collection for evaluating the retrieval performance of service discovery approaches (matchmakers, search engines) for the standard SAWSDL and WSML, while there are two publicly available for OWL-S (OWLS-TC2, SWS-TC). There are no solutions for the integrated matching of different services that are specified in different languages like SAWSDL, OWL-S and WSML. Relevant work on refactoring OWL-S and WSML to the standard SAWSDL is ongoing.

Some major challenges of research and development in the domain of Semantic Web service composition planning are as follows.

- Scalable and resource efficient approaches to service composition planning under uncertainty and their use in real-world applications of the Web 3.0 and in intelligent pervasive service applications of the so called “Internet of Things” that is envisioned to interlink all kinds of computing devices without limit on the global scale.
- Efficient means of distributed composition planning of Semantic Web Services in peer-to-peer and grid computing environments.
- Easy to use tools for the common user to support discovery, negotiation, composition and execution Semantic Web Services in one framework for different Semantic Web Service formats like the standard SAWSDL, and non-standards like OWL-S, WSML, and SWSL.
- Interleaving of service composition planning with negotiation in competitive settings.

4.6 Summary

This chapter provided a brief romp through the fields of Semantic Web service discovery and composition planning. We classified existing approaches, discussed representative examples and commented on the interrelationships between both service coordination activities. Despite fast paced research and development in the past years world wide, Semantic Web service technology still is commonly considered immature with many open theoretical and practical problems as mentioned above. However, its current convergence with Web 2.0 towards a so-called service Web 3.0 in an envisioned Internet of Things holds promise to effectively revolutionize computing applications for our everyday life.

References

- [1] K. Aberer, P. Cudre-Mauroux, M. Hauswirth: Semantic Gossiping: fostering semantic interoperability in peer data management systems. S. Staab, H. Stuckenschmidt (eds.): *Semantic Web and Peer-to-Peer*, Springer, Chapter 13, 2006.
- [2] G. Alonso, F. Casati, H. Kuno, V. Machiraju: *Web Services*. Springer, 2003
- [3] S. Agarwal, S. Handschuh, S. Staab: Annotation, composition and invocation of Semantic Web Services. *Web Semantics*, 2, 2004.
- [4] G. Agre, Z. Marinova: An INFRAWEBs Approach to Dynamic Composition of Semantic Web Services. *Cybernetics and Information Technologies (CIT)*, 7(1), 2007.
- [5] M.S. Aktas, G. Fox, M. Pierce: Managing Dynamic Metadata as Context. Proceedings of Intl. Conference on Computational Science and Engineering (ICCSE), Istanbul, 2005
- [6] S. Amer-Yahia, C. Botev, J. Shanmugasundaram: TeXQuery: A Full-Text Search Extension to XQuery. Proceedings of the World-Wide-Web Conference WWW 2004, 2004.
- [7] A. Ankolekar, M. Paolucci, K. Sycara: Spinning the OWL-S Process Model - Toward the Verification of the OWL-S Process Models. Proceedings of International Semantic Web Services Workshop (SWSW), 2004
- [8] I.B. Arpinar, B. Aleman-Meza, R. Zhang, A. Maduko: Ontology-driven Web Services composition platform. Proc. of IEEE International Conference on E-Commerce Technology CEC, San Diego, USA, IEEE Press, 2004.
- [9] M.A. Aslam, S. Auer, J. Shen: From BPEL4WS Process Model to Full OWL-S Ontology. Proceedings of 2nd European Conference on Semantic Web Services ESWC, Buda, Montenegro, 2006.

- [10] F. Baader, C. Lutz, M. Milicic, U. Sattler, F. Wolter: Integrating Description Logics and action formalisms: First results. Proc. 20th National Conference on Artificial Intelligence (AAAI), Pittsburgh, USA, AAAI Press, 2005
- [11] J. Bae, L. Liu, J. Caverlee, W.B. Rouse: Process Mining, Discovery, and Integration using Distance Measures. Proceedings of International Conference on Web Services ICWS, 2006.
- [12] S. Bansal, J. Vidal: Matchmaking of Web Services Based on the DAMLS Service Model. Proc. International Joint Conference on Autonomous Agents and Multiagent Systems AAMAS, 2003.
- [13] U. Basters and M. Klusch: RS2D: Fast Adaptive Search for Semantic Web Services in Unstructured P2P Networks. Proceedings 5th Intl. Semantic Web Conference (ISWC), Athens, USA, Lecture Notes in Computer Science (LNCS), 4273:87-100, Springer, 2006.
- [14] A. Bernstein, C. Kiefer: Imprecise RDQL: Towards Generic Retrieval in Ontologies Using Similarity Joins. Proceedings ACM Symposium on Applied Computing, Dijon, France, ACM Press, 2006.
- [15] P. Bertoli, A. Cimatti, P. Traverso: Interleaving Execution and Planning for Nondeterministic, Partially Observable Domains. Proceedings of European Conference on Artificial Intelligence (ECAI), 2004.
- [16] D. Bianchini, V. De Antonellis, M. Melchiori, D. Salvi: Semantic-enriched Service Discovery. Proceedings of IEEE ICDE 2nd International Workshop on Challenges in Web Information Retrieval and Integration (WIRI06), Atlanta, Georgia, USA, 2006.
- [17] W. Binder, I. Constantinescu, B. Faltings, K. Haller, C. Tuerker: A Multi-Agent System for the Reliable Execution of Automatically Composed Ad-hoc Processes. Proceedings of the 2nd European Workshop on Multi-Agent Systems (EUMAS), Barcelona, Spain, 2004.
- [18] L. Botelho, A. Fernandez, B. Fries, M. Klusch, L. Pereira, T. Santos, P. Pais, M. Vasirani: Service Discovery. In M. Schumacher, H. Helin (Eds.): CASCOM - Intelligent Service Coordination in the Semantic Web. Chapter 10. Birkh"auser Verlag, Springer, 2008.
- [19] C. Caceres, A. Fernandez, H. Helin, O. Keller, M. Klusch: Context-aware Service Coordination for Mobile Users. Proceedings IST eHealth Conference, 2006.
- [20] F. Casati, M.C. Shan: Dynamic and Adaptive Composition of E-services. *Information Systems*, 6(3), 2001.
- [21] CASCOM Project Deliverable D3.2: Conceptual Architecture Design. September 2005. www.ist-cascom.org

- [22] D. Chakraborty, F. Perich, S. Avancha, A. Joshi: DReggie: Semantic Service Discovery for M-Commerce Applications. Proceedings of the International Workshop on Reliable and Secure Applications in Mobile Environment, 2001.
- [23] H. Chen, A. Joshi, and T. Finin: Dynamic service discovery for mobile computing: Intelligent agents meet JINI in the aether. 4(4):343-354, 2001.
- [24] S. Colucci, T.C. Di Noia, E. Di Sciascio, F.M. Donini, M. Mongiello: Concept Abduction and Contraction for Semantic-based Discovery of Matches and Negotiation Spaces in an E-Marketplace. *Electronic Commerce Research and Applications*, 4(4):345-361, 2005.
- [25] I. Constantinescu, B. Faltings: Efficient matchmaking and directory services Proceedings of IEEE Conference on Web Intelligence WI, 2003.
- [26] R. Dearden, N. Meuleauy, S. Ramakrishnany, D.E. Smith, R. Washington: Incremental Contingency Planning. Proc. of ICAPS-03 Workshop on Planning under Uncertainty, Trento, Italy, 2003.
- [27] E. Della Valle, D. Cerizza, I. Celino: The Mediators Centric Approach to Automatic Web Service Discovery of Glue. Proceedings of 1st International Workshop on Mediation in Semantic Web Services (MEDIATE), CEUR Workshop proceedings, 168, 2005.
- [28] G. Denker, L. Kagal, T. Finin, M. Paolucci, K. Sycara: Security For DAML Web Services: Annotation and Matchmaking. Proceedings of the Second International Semantic Web Conference (ISWC 2003), USA, 2003.
- [29] T. Di Noia, E.D. Sciascio, F.M. Donini, M. Mogiello: A System for Principled Matchmaking in an Electronic Marketplace. *Electronic Commerce*, 2004.
- [30] T. Di Noia, E. Di Sciascio, F.M. Donini: Semantic Matchmaking as Non-Monotonic Reasoning: A Description Logic Approach. *Artificial Intelligence Research (JAIR)*, 29:269-307, 2007.
- [31] J. Domingue, S. Galizia, L. Cabral: Choreography in IRS-III: Coping with Heterogeneous Interaction Patterns in Web Services. Proc. International Semantic Web Conference, LNAI, Springer, 2005.
- [32] D. Fahland, W. Reisig: ASM-based semantics for BPEL: The negative Control Flow. Proceedings of the 12th International Workshop on Abstract State Machines (ASM'05), 2005.
- [33] D. Fensel, F. van Harmelen: Unifying reasoning and search to Web scale. *IEEE Internet Computing*, March/April 2007.
- [34] A. Fernandez, M. Vasirani, C. Caceres, S. Ossowski: A role-based support mechanism for service description and discovery. In: huang et al. (eds.), Service-Oriented Computing: Agents, Semantics, and Engineering. LNCS 4504, Springer, 2006.

- [35] U. Furbach, M. Maron, K. Read: Location based informationsystems. *Künstliche Intelligenz*, 3/07, BöttcherIT, 2007.
- [36] M. Ghallab, D. Nau, P. Traverso: *Automated planning*. Elsevier, 2004.
- [37] S. Grimm: Discovery - Identifying relevant services. In [104], 2007.
- [38] S. Grimm, B. Motik, C. Preist: Matching semantic service descriptions with local closed-world reasoning. *Proc. European Semantic Web Conference (ESWC)*, Springer, LNCS, 2006.
- [39] L. Guo, F. Shao, C. Botev, J. Shanmugasundaram: XRANK: Ranked Keyword Search over XML Documents. *Proceedings of the 2003 ACM SIGMOD International Conference on Management of Data*, San Diego, USA, 2003.
- [40] P. Haase, R. Siebes, F. van Harmelen: Expertise-based Peer selection in Peer-to-Peer Networks. *Knowledge and Information Systems*, Springer, 2006
- [41] L. Henoque, M. Kleiner: Composition - Combining Web Service Functionality in Composite Orchestrations. Chapter 9 in [104], 2007.
- [42] D. Hull, U. Sattler, E. Zolin, R. Stevens, A. Bovykin, I. Horrocks: Deciding semantic matching of stateless services. *Proc. 21st National Conference on Artificial Intelligence (AAAI)*, AAAI Press, 2006
- [43] D. Hutter, M. Klusch, M. Volkamer: Information Flow Analysis Based Security Checking of Health Service Composition Plans. *Proceedings of the 1st European Conference on eHealth*, Fribourg, Switzerland, 2006.
- [44] D. Hutter, M. Volkamer, M. Klusch, A. Gerber: Provably Secure Execution of Composed Semantic Web Services. *Proceedings of the 1st International Workshop on Privacy and Security in Agent-based Collaborative Environments (PSACE 2006)*, Hakodate, Japan, 2006.
- [45] M.C. Jäger, G. Rojec-Goldmann, C. Liebetrueth, G. Mühl, K. Geihs: Ranked Matching for Service Descriptions Using OWL-S. *Proceedings of 14. GI/VDE Fachtagung Kommunikation in Verteilten Systemen KiVS*, Kaiserslautern, 2005
- [46] L. Kagal, T. Finin, M. Paolucci, N. Srinivasan, K. Sycara, G. Denker: Authorization and Privacy for Semantic Web Services. *IEEE Intelligent Systems*, July/August, 2004.
- [47] F. B. Kashani, C.C. Shen, C. Shahabi: SWPDS: Web Service peer-to-peer discovery service. *Proceedings of Intl. Conference on Internet Computing*, 2004.
- [48] U. Keller, R. Lara, H. Lausen, A. Polleres, D. Fensel: Automatic Location of Services. *Proceedings of the 2nd European Semantic Web Conference (ESWC)*, Heraklion, Crete, LNCS 3532, Springer, 2005.

- [49] F. Kaufer and M. Klusch: Hybrid Semantic Web Service Matching with WSMO-MX. Proc. 4th IEEE European Conference on Web Services (ECOWS), Zurich, Switzerland, IEEE CS Press, 2006
- [50] F. Kaufer and M. Klusch: Performance of Hybrid WSML Service Matching with WSMO-MX: Preliminary Results. Proc. First Intl. Joint ISWC Workshop SMR2 2007 on Service Matchmaking and Resource Retrieval in the Semantic Web, Busan, Korea, 2007.
- [51] C. Kiefer, A. Bernstein: The Creation and Evaluation of iSPARQL Strategies for Matchmaking. Proceedings of European Semantic Web Conference, Springer, 2008.
- [52] T. Kleemann, A. Sinner: Description logic based matchmaking on mobile devices. Proceedings of 1st Workshop on Knowledge Engineering and Software Engineering (KESE 2005), 2005.
- [53] M. Klein, B. König-Ries: Coupled Signature and Specification Matching for Automatic Service Binding. European Conference on Web Services (ECOWS 2004), Erfurt, 2004.
- [54] M. Klusch, B. Fries: Hybrid OWL-S Service Retrieval with OWLS-MX: Benefits and Pitfalls. Proceedings 1st International Joint Workshop on Service Matchmaking and Resource Retrieval in the Semantic Web (SMR2), Busan, Korea, CEUR vol. 243, 2007.
- [55] M. Klusch, P. Kapahnke, B. Fries: Hybrid Semantic Web Service Retrieval: A Case Study With OWLS-MX. Proceedings of 2nd IEEE International Conference on Semantic Computing (ICSC), Santa Clara, USA, IEEE Press, 2008.
- [56] M. Klusch, A. Gerber, M. Schmidt: Semantic Web Service Composition Planning with OWLS-XPlan. Proc. 1st Intl. AAAI Fall Symposium on Agents and the Semantic Web, Arlington VA, USA, AAAI Press, 2005.
- [57] M. Klusch, K-U. Renner: Dynamic Re-Planning of Composite OWL-S Services. Proc. 1st IEEE Workshop on Semantic Web Service Composition, Hongkong, China, IEEE CS Press, 2006.
- [58] M. Klusch, K. Sycara: Brokering and Matchmaking for Coordination of Agent Societies: A Survey. In: Coordination of Internet Agents, A. Omicini et al. (eds.), Springer
- [59] M. Klusch, B. Fries, K. Sycara: Automated Semantic Web Service Discovery with OWLS-MX. Proc. 5th Intl. Conference on Autonomous Agents and Multi-Agent Systems (AAMAS), Hakodate, Japan, ACM Press, 2006
- [60] P. Küngas, M. Matskin: Semantic Web Service Composition through a P2P-Based Multi-Agent Environment. Proc. of the Fourth International Workshop on Agents and Peer-to-Peer Computing (in conjunction with AAMAS 2005), Utrecht, Netherlands, LNCS 4118, 2006.

- [61] U. Küster, B. König-Ries, M. Stern, M. Klein: DIANE: An Integrated Approach to Automated Service Discovery, Matchmaking and Composition. Proceedings of the World Wide Web Conference WWW, Banff, Canada, ACM Press, 2007.
- [62] S. Lamparter, A. Ankolekar: Automated Selection of Configurable Web Services. 8. Internationale Tagung Wirtschaftsinformatik. Universitätsverlag Karlsruhe, Karlsruhe, Germany, March 2007.
- [63] F. Lecue, A. Leger: Semantic Web Service composition through a match-making of domain. Proc. of 4th IEEE European Conference on Web Services (ECWS), Zurich, 2006.
- [64] F. Lecue, A. Delteil, A. Leger: Applying Abduction in Semantic Web Service Composition. Proceedings of IEEE International Conference on Web Services (ICWS 2007), 2007.
- [65] L. Li, I. Horrocks: A software framework for matchmaking based on semantic Web technology. Proceedings of the world wide Web conference (WWW), Budapest, 2003.
- [66] J. Liu, H. Zhuge: A Semantic-Link-Based Infrastructure for Web Service. Proc. of the International World Wide Web Conference, 2005.
- [67] S. Liu, P. Küngas, M. Matskin: Agent-Based Web Service Composition with JADE and JXTA. Proc. of Intl Conference on Semantic Web and Web Services (SWWS), Las Vegas, USA, 2006.
- [68] A. Löser, C. Tempich, B. Quilitz, W.-T. Balke, S. Staab, W. Nejdl: Searching Dynamic Communities with Personal Indexes. Proceedings of Internatioanal Semantic Web Conference, 2005.
- [69] N. Lohmann: A Feature-Complete Petri Net Semantics for WS-BPEL 2.0. Proceedings of the Workshop on Formal Approaches to Business Processes and Web Services (FABPWS'07), 2007.
- [70] B.T. Loo, R. Huebsch, I. Stoica, J.M. Hellerstein: The Case for a Hybrid P2P Search Infrastructure. Proceedings of rd Intl Workshop on P2P Systems (IPTPS), USA, Springer, LNCS, 2004.
- [71] Q. Lu, P. Cao, E. Cohen, K. Li, S. Shenker: Search and Replication in Unstructured Peer-to-Peer Networks. Proceedings of ACM 6th ACM International Conference on Supercomputing ICS, New York, USA, 2002.
- [72] A. Martens: Analyzing Web Service based Business Processes. Proceedings of Workshop on Fundamental Approaches to Software Engineering FASE, 2005.
- [73] S. McIlraith, T.C. Son: Adapting Golog for composition of Semantic Web Services. Proc. International Conference on Knowledge Representation and Reasoning KRR, Toulouse, France, 2002.

- [74] S. Narayanan, S. McIlraith: Simulation, verification and automated composition of Web Services. Proc. of 11th International Conference on the World Wide Web (WWW), Hawaii, 2002.
- [75] B. Medjahed, A. Bouguettyaya, A.K. Elmagarmid: Composing Web Services on the semantic Web. *Very Large Data Bases (VLDB)*, 12(4), 2003.
- [76] B.T. Messmer: New approaches on graph matching. PhD Thesis, University of Bern, Switzerland, 1995.
- [77] M. Milicic: Planning in Action Formalisms based on DLS: First Results. Proceedings of the Intl Workshop on Description Logics, 2007.
- [78] D.S. Milojevic, V. Kalogeraki, R. Lukose, K. Nagaraja, J. Pruyne, B. Richard, S. Rollins, Z. Xu: Peer-to-peer computing. Technical Report HPL-2002-57, Hewlett-Packard, 2002.
- [79] T. Möller, H. Schuldt, A. Gerber, M. Klusch: Next Generation Applications in Healthcare Digital Libraries using Semantic Service Composition and Coordination. *Health Informatics*, 12 (2):107-119, SAGE publications, 2006.
- [80] I. Müller, R. Kowalczyk, P. Braun: Towards Agent-Based Coalition Formation for Service Composition. Proceedings of the IEEE International Conference on Intelligent Agent Technology, Washington, USA, 2006.
- [81] M. Paolucci, T. Kawamura, T.R. Payne, K. Sycara: Semantic Matching of Web Services Capabilities. Proceedings of the 1st International Semantic Web Conference (ISWC2002), 2002.
- [82] M. Paolucci, K. Sycara, T. Nishimara, N. Srinivasan: Using DAML-S for P2P Discovery. Proc. of International Conference on Web Services, Erfurt, Germany, 2003.
- [83] M. Papazoglou: Web Services: Principles and Technology. Pearson - Prentice Hall, September 2007.
- [84] J. Peer: Web Service Composition as AI Planning: A Survey. Technical Report, University of St. Gallen, Switzerland, 2005. Available at elektra.mcm.unisg.ch/pbwsc/docs/pfwsc.pdf
- [85] J. Peer: A POP-Based Replanning Agent for Automatic Web Service Composition. Proceedings of the 2nd European Semantic Web Conference (ESWC), Heraklion, Crete, LNCS 3532, Springer, 2005.
- [86] A. Pfalzgraf: Ein robustes System zur automatischen Komposition semantischer Web Services in SmartWeb. Master Thesis, University of the Saarland, Saarbrücken, Germany, Juni 2006.
- [87] M. Pistore, P. Traverso: Planning as model checking for extended goals in non-deterministic domains. In: Proceedings of the 7th International Joint Conference on Artificial Intelligence (IJCAI-01), 2001.

- [88] M. Pistore, P. Roberti, P. Traverso: Process-Level Composition of Executable Web Services: On-the-fly Versus Once-for-all Composition Proceedings of the 2nd European Semantic Web Conference (ESWC), Heraklion, Crete, LNCS 3532, Springer, 2005.
- [89] M. Pistore, P. Traverso, P. Bertoli, A. Marconi: Automated synthesis of composite BPEL4WS Web Services. Proceedings of the 2005 IEEE International Conference on Web Services, Orlando, USA, IEEE Press, 2005.
- [90] C. Preist: Semantic Web Services - Goals and Vision. Chapter 6 in [104], 2007.
- [91] C. Preist, C. Bartolini, A. Byde: Agent-based service composition through simultaneous negotiation in forward and reverse auctions. Proceedings of the 4th ACM Conference on Electronic Commerce, San Diego, California, USA, 2003.
- [92] J. Rao, P. Kuengas, M. Matskin: Composition of Semantic Web Services using Linear Logic theorem proving. *Information Systems*, 31, 2006.
- [93] K.-U. Renner, P. Kapahnke, B. Blankenburg, M. Klusch: OWLS-XPlan 2.0 - A Dynamic OWL-S Service Composition Planner. BMB+F project SCALLOPS, Internal Project Report, DFKI Saarbrücken, Germany, 2007. www.dfki.de/~klusch/owls-xplan2-report-2007.pdf
- [94] A. Rosenfeld, C. Goldman, G. Kaminka, S. Kraus: An Agent Architecture for Hybrid P2P Free-Text Search. Proceedings of 11th Intl Workshop on COoperative Information Agents (CIA), Delft, Springer, LNAI 4676, 2007.
- [95] M. Schlosser, M. Sintek, S. Decker, W. Nejdl: A Scalable and Ontology-based P2P Infrastructure for Semantic Web Services. Proceedings of 2nd IEEE Intl Conference on Peer-to-Peer Computing (P2P), Linkoping, Sweden, 2003
- [96] B. Schnizler, D. Neumann, D. Veit, C. Weinhardt: Trading Grid Services - A Multi-attribute Combinatorial Approach. *European Journal of Operational Research*, 2006.
- [97] E. Sirin, J. Hendler, B. Parsia: Semi-automatic Composition of Web Services using Semantic Descriptions. Proceedings of Intl Workshop on Web Services: Modeling, Architecture and Infrastructure workshop in conjunction with ICEIS conference, 2002.
- [98] E. Sirin, B. Parsia, J. Hendler: Filtering and Selecting Semantic Web Services with Interactive Composition Techniques. *IEEE Intelligent Systems*, July/August, 2004.
- [99] E. Sirin, B. Parsia, D. Wu, J. Hendler, D. Nau: HTN planning for Web Service composition using SHOP2. *Web Semantics*, 1(4), Elsevier, 2004.
- [100] D.E. Smith, D.S. Weld: Conformant Graphplan. Proc. of 15th AAAI Conference on AI, Pittsburgh, USA, 1998.

- [101] B. Srivastava, J. Koehler: Web Service Composition: Current Solutions and Open Problems. Proceedings of the ICAPS 2003 Workshop on Planning for Web Services, 2003.
- [102] S. Staab, H. Stuckenschmidt (eds.): Semantic Web and Peer-to-Peer. Springer, 2006.
- [103] M. Stollberg, U. Keller, H. Lausen, S. Heymans: Two-phase Web Service discovery based on rich functional descriptions. Proceedings of European Semantic Web Conference, Buda, Montenegro, LNCS, Springer, 2007.
- [104] R. Studer, S. Grimm, A. Abecker (eds.): Semantic Web Services. Concepts, Technologies, and Applications. Springer, 2007.
- [105] K. Sycara, M. Klusch, S. Widoff, J. Lu: LARKS: Dynamic Matchmaking Among Heterogeneous Software Agents in Cyberspace. *Autonomous Agents and Multi-Agent Systems*, 5(2):173 - 204, Kluwer Academic, 2002.
- [106] D. Trastour, C. Bartolini, C. Priest: Semantic Web Support for the Business-to-Business E-Commerce Lifecycle. Proceedings of the International World Wide Web Conference (WWW), 2002.
- [107] P. Traverso, M. Pistore: Automated Composition of Semantic Web Services into Executable Processes. Int Semantic Web Conference, LNCS 3298, Springer, 2004.
- [108] D. Tsoumakos, N. Roussopoulos: Adaptive Probabilistic Search (APS) for Peer-to-Peer Networks. Proc. Int. IEEE Conference on P2P Computing, 2003.
- [109] R. Vaculin, K. Sycara: Towards automatic mediation of OWL-S process models. IEEE International Conference on Web Services (ICWS 2007), 2007.
- [110] W.M.P. van der Aalst, A.J.M.M. Weijters: Process mining: a research agenda. *Computers in Industry*, 53, 2004.
- [111] K. Verma, K. Sivashanmugam, A. Sheth, A. Patil, S. Oundhakar, J. Miller: METEORS WSDI: A Scalable P2P Infrastructure of Registries for Semantic Publication and Discovery of Web Services. *Information Technology and Management*, Special Issue on Universal Global Integration, Vol. 6, No. 1, 2005.
- [112] L.H. Vu, M. Hauswirth, F. Porto, K. Aberer: A Search Engine for QoS-enabled Discovery of Semantic Web Services. Ecole Polytechnique Federal de Lausanne, LSIR-REPORT-2006-002, Switzerland, 2006. Also available in the Special Issue of the International Journal of Business Process Integration and Management (IJBPIIM) (2006).
- [113] Z. Wu, K. Gomadam, A. Ranabahu, A. Sheth, J. Miller: Automatic Composition of Semantic Web Services using Process Mediation. Proceedings of the 9th Intl. Conf. on Enterprise Information Systems ICES 2007, Funchal, Portugal, 2007.

- [114] J. Yan, R. Kowalczyk, J. Lin, M.B. Chhetri, S.K.Goh, J. Zhang: Autonomous service level agreement negotiation for service composition provision. *Future Generation Computing Systems*, 23(6), Elsevier, 2007.
- [115] A.M. Zaremski, J.M. Wing: Specification Matching of Software Components. *ACM Transactions on Software Engineering and Methodology*, 6(4), 1997.