

A Flexible and Semantic-aware Publication Infrastructure for Web Services

Luciano Baresi, Matteo Miraz, and Pierluigi Plebani

Dipartimento di Elettronica e Informazione – Politecnico di Milano
Piazza Leonardo da Vinci 32, 20133 Milano (Italy)
{baresi, miraz, plebani}@elet.polimi.it

Abstract. This paper presents an innovative approach for the publication and discovery of Web services. The proposal is based on two previous works: DIRE (DIstributed REgistry), for the user-centered distributed replication of service-related information, and URBE (UDDI Registry By Example), for the semantic-aware match making between requests and available services. The integrated view also exploits USQL (Unified Service Query Language) to provide users with higher level and homogeneous means to interact with the different registries. The proposal improves background technology in different ways: we use and integrate USQL as a high-level language to state service requests, widen user notifications based on URBE semantic matching, and apply URBE match making to all the facets with which services can be described in DIRE. All these new concepts are demonstrated on a simple example scenario.

1 Introduction

So far, the publication and discovery of Web services [1] have been tackled in several different ways without identifying any widely accepted solutions. While the community agrees on WSDL (Web Service Description Language) for the syntactical description of Web services' functionality, and on WS-BPEL (Business Process Execution Language, [2]) as means to compose existing services, the efficient and effective *exposition* and *retrieval* of services are still open problems and have only found partial solutions so far. UDDI (Universal Description Discovery Integration, [3]) and ebXML Registry [4] are proposals for registries whose actual use is limited. The semantic community has been pushing richer service descriptions [5–7], based on description logics and ontologies, to support accurate service discovery. METEOR-S [8] and Pyramid-S [9] combine the previous proposals and offer semantically-enriched ways for the service publication.

The lack of a winning solution pushed us to further analyze the problem and concentrate on the distributed publication of services as a way to improve both their exposition and their retrieval. This paper proposes a novel approach for the publication of Web services to ease their discovery. More specifically, we think that the information about available services must be moved closer to their possible users, and it must be organized in a user-centric way. Centralized general-purpose repositories (e.g, the UDDI Business Registry started by IBM,

Microsoft, and SAP) have already failed¹, and even if all the main proposals have moved towards distributed approaches, we think that this distribution cannot be super-imposed. Our proposal is to let the users fully control their registries (i) by defining what they want to share with the others and (ii) by specifying the services potentially available on external registries they are interested in. Thus, we concentrate on the distributed *user-centered* propagation of service information and on the discovery features that such a distribution enables. The discovery, in particular, is enriched with the adoption of semantic-based analysis to improve the responsiveness of the system and help the users with solutions (services) that are close enough to what they would have liked to get (even if they do not fully match their expectations).

The proposed interaction among registries exploits a *publish and subscribe* [10] (P/S, hereafter) communication infrastructure to allow for flexible and dynamic interactions among registries. Unlike the current available registries, this means that each registry can decide the services it wants to publish, that is, the services it wants to share with the others. Similarly, it can declare its interests by means of special-purpose subscriptions. The infrastructure ensures that as soon as a registry publishes the information about one of its services, this same information is propagated to (and replicated on) all the registries that had declared their interest. Subscriptions (and unsubscriptions) can be issued dynamically and thus each registry can accommodate and tailor its interests (i.e., those of their users) while in operation.

The second key message of the paper is that oftentimes users are not only interested in services that fully and exactly match their requests, but they also want to know if there are “similar” solutions, that is, services that suitably adapted can be used instead of the ones part of the original request. This requirement is tackled in the paper in two different and orthogonal ways. User requests are formulated in a technology neutral and high-level query language, called USQL (Unified Service Query Language, [11]), and are then automatically translated into subscriptions suitably distributed through the communication infrastructure. On the other hand, the dispatching is powered with matchmaking capabilities to provide the different registries with semantically-enriched notifications, that is, information about services whose match with the original request (subscription) is within a given threshold.

The work presented in this paper builds on top of two existing proposals: namely DIRE (DIstributed REgistry, [12]), as for the communication framework among registries, and URBE (UDDI Registry By Example, [13]) as for the matchmaking and semantic awareness. The integration of the two proposals allows us to consider a semantically-enabled replication infrastructure that supports different registry technologies (UDDI, ebXML, and the SeCSE registry²) by means of JAXR (Java API for XML Registries, [14]), and also the *facet*-based [15] description of services, that is, each particular aspect (XML document) of a ser-

¹ <http://webservices.sys-con.com/read/164624.htm>

² <http://secse.eng.it>

vice is described by means of a proper aspect, be it specific or shared with other services.

Besides the obvious integration of the two proposals, the novel contributions of this paper lie in: (i) the use of USQL as high-level language to state service requests, along with its automatic translation in terms of subscriptions for the communication infrastructure, (ii) the widening of notifications based on URBE semantic matching, and (iii) the extension of the URBE matching to all the facets with which services can be described.

The rest of the paper is organized as follows. Section 2 introduces an example scenario to motivate the proposal presented in the paper, while Section 3 summarizes background technologies. Section 4 describes the complete infrastructure introduced here along with the new features. Section 5 surveys some related proposals and Section 6 concludes the paper.

2 Example scenario

Even if the UDDI Business Registries by IBM, Microsoft, and SAP are not operated anymore, alternative “global” Web service registries are still available, and among the others XMethods³ and Wsoogle⁴ are currently used by many providers to publish their services. These registries are used worldwide, host Web services of any kind, and provide facilities to ease the discovery. Since the number of available services is always increasing, the approach introduced in this paper can be seen as a means to further simplify and increase the effectiveness of service discovery.

Figure 1 introduces our running example, where besides the two “general-purpose” registries introduced above, we have three other proprietary registries that host Web services of different application domains. Notice that this is nothing but a way to exemplify how our approach works; further technical considerations behind this example are outside the scope of the paper.

The first registry belongs to a company specialized in software development for healthcare solutions. They are interested in Web services able to support as many activities as possible in this application domain. The second is run by a tour operator willing to improve its Web site with mash-up services, while the third is operated by a community of chess players who want to be aware of new opportunities (Web services) to play chess over the Internet. At this stage, the managers of these registries are in charge of periodically browsing the general-purpose ones to find the services of interest and update their local copies.

Each requester is interested in services of different categories, but also the way in which they are able to express their requirements is different and depends on their skills. For instance, on the average, chess players do not know WSDL, but they can easily express their requirements using chess- and QoS-related keywords (e.g., *chess* or *free chess server*). In contrast, the software company might look for Web services with particular WSDL interfaces. In addition, the

³ <http://www.xmethods.net>

⁴ <http://www.wsoogle.com>

software company might also become service provider and offer Web services that help citizens to locate hospitals in a given area. In this case, instead of publishing the new Web service on one of the main registries, the software house would like to publish it locally and preferably on a single server.

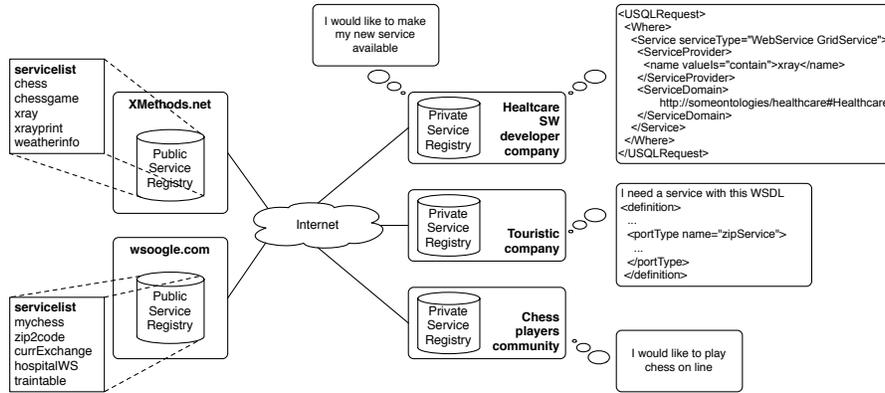


Fig. 1. Example scenario.

All these different users might browse XMethods and Wsoogle to find what they need, but this activity would very time consuming and the actual results would heavily depend on the ability of who works on service discovery. They would also love to have their own registries with only the Web services that might be of interest, and with no need for period updates.

This is exactly what our solution proposes: each registry connected to the communication infrastructure described in Section 4 only has to declare its interests in the way its users prefer and then the infrastructure grabs relevant services directly from where they are published (mainly the two big repositories, in our example). When one of the users also holds the role of service provider, the infrastructure automatically publishes the new service onto the other interested registries. Since Web services can be published and unpublished, the infrastructure is also in charge of updating the proprietary replicas as soon as new information (services) becomes available. In this scenario, all the services published in the general purpose registries are public by definition. In the same way, these registries are interested in all the services in the private registries that are declared as public by their providers.

3 Background

This section briefly recalls the main concepts behind the technologies used in the paper. Given the already identified connections with DIRE, URBE, and USQL, we think it is better to sketch their main features here to provide the reader with

a self-contained paper, and also to highlight those elements that will be used in the next sections.

3.1 DIRE

DIRE (DIstributed REgistry, [12])⁵ provides a common service model for heterogeneous registries and makes them communicate by means of a P/S communication bus.

DIRE is in line with those approaches that tend to unify the service model (e.g., JAXR and USQL). Business data are rendered by means of pre-defined elements called **Organizations**, **Services**, and **ServiceBindings**, with the meaning that these elements usually assume in existing registries. Technical data are described by means of typed **Facets**, where each facet addresses a particular feature of the service by using an XML language. **StandardFacets** characterize recurring features (for example the compliance with an abstract interface) and we assume that they are shared among different services. **SpecificFacets** describe the peculiarities of the different services (for example, particular SLAs or additional technical information). Users can attach new facets to services, even if they are not the provider, to customize the way services are perceived by the different registries (users) and to let them share this information with the other components attached to the communication bus.

The peculiarity of the communication bus, which is based on a distributed P/S middleware called *ReDS* [16], is that components do not interact directly, but through a *dispatcher*. Each component can *publish* its messages on the dispatcher and decides the messages it wants to listen to (*subscribe/unsubscribe*). The dispatcher forwards (*notifies*) received messages to all registered components. This means that the different registries must declare their interests for particular services. This is done by means of special-purpose filters that refer to shared standard facets or embed XPath expressions on specific facets.

Each publication —on a particular registry connected to the bus— is propagated to those registries that subscribed to it. Once the information is received, each registry can either discard it, or store it locally. The goal is to disseminate the information about services based on interests and requests, instead of according to predefined semantic similarities. This way, DIRE supports a two-phase discovery process. The registry retrieves the services the user is interested in from the P/S infrastructure. Users always search for the services they need locally.

A *delivery manager* is attached to each registry and acts as *facade*, that is, it is the intermediary between the registry and the bus and manages the information flow in the two directions. The delivery manager is responsible for the policies for both sharing the services in the registry and selecting the services that must be retrieved from the P/S infrastructure. Its adoption does not require modifications to the publication and discovery processes used by the different users. They keep interacting with the repository they were used to, but published

⁵ <http://code.google.com/p/delivery-manager/>

services are distributed through the P/S infrastructure, which in turn provides information about the services published by the others (if they are of interest).

Notice that a registry can connect to the bus and declare its interests at any time. The infrastructure guarantees that a registry can always retrieve the information it is interested in by means of *lease* contracts. The lease period, which is configurable at run-time, guarantees that the information about services is re-transmitted periodically. It is also the maximum delay with which a registry is notified about a service. Moreover, if the description of a service changes, the lease guarantees that the new data are distributed to all subscribed registries within the period.

3.2 URBE

URBE⁶ (UDDI Registry By Example) is an extension to typical UDDI Registries to support content-based queries, i.e. users can retrieve services whose operations have a given input or output. URBE's users submit a WSDL (or a SAWSDL, Semantic Annotated WSDL, [7]) description to express how the signature of requested Web services should be. As result, URBE returns a ranked list of Web services whose signature is similar to the submitted one. URBE supports service substitutability at both design- and run-time: i.e., it finds a substitute Web service that exposes an interface equal to or at least similar to the interface of the substituted service. In addition, URBE supports the top-down design of WS-BPEL processes. Traditional design approaches push the designer to start identifying the potential partners and then design the WS-BPEL process by exploiting the previously selected WSDL interfaces (bottom-up approach). URBE allows the designer to initially focus on the process definition and then discover the Web services able to perform required invocations.

URBE's *similarity engine* compares WSDL descriptions of Web services. Assuming that users express their queries using WSDL, this component compares the submitted WSDL with the WSDL of all the Web services in the registry. Each comparison relies on function $WSDLsim : (wsdl_q, wsdl_p) \rightarrow [0..1]$, where the higher the result is, the higher the similarity between the two Web services is [17]. This value is obtained recursively by analyzing the overall signature, the operations, and their parameters; it reflects the amount of requirements expressed by query Q satisfied by the signature of Web service P . More in detail, for each operation in Q , the similarity engine finds the operation in P with maximum similarity. This similarity depends on the similarity between the operations' names (calculated by $opSim$) and the similarity of their input and output parameters (calculated by $parSim$). Finally, the similarity of parameters depends on the similarity of the parameters' names and their data types. A high level overview of the similarity evaluation process is given in Figure 2.

As a consequence, the similarity between two signatures heavily depends on the names assigned to the whole service, available operations, and exchanged parameters. The comparison between terms relies on a term similarity function

⁶ <http://black.elet.polimi.it/urbe>

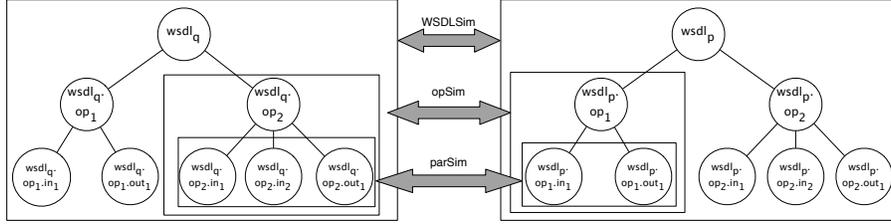


Fig. 2. Example similarity evaluation.

$termSim : (t_i, t_j) \rightarrow [0..1]$. This function returns a value that reflects how the two terms t_i and t_j are semantically close: 1 if t_i and t_j are synonym, 0 if they are antonym. To achieve this goal, $termSim$ relies on two main ontologies: a *domain specific ontology* and a *general purpose ontology*. The first includes terms related to a given application domain. We assume that this ontology can be built by domain experts by analyzing the terms included in the Web services published in the registry. The latter includes all the possible terms (at this stage we adopt Wordnet⁷). We decided to rely on both ontologies since the domain specific ontology offers more accuracy in the relationships among terms, while the general purpose one offers wider coverage. This happens because in a general purpose ontology a word may have more than one *synset*, each corresponding to a different meaning. In contrast, we assume that in a domain specific ontology each word has a unique meaning with respect to the domain itself.

Name similarity depends on the way two names are connected inside the ontology [18]. If we assume that WSDL descriptions are generated automatically, for example from Java classes, it is possible that the names of operations and parameters reflect the naming convention usually adopted by programmers: `getData` or `currencyExchange` are more frequent than the simple names directly included in the ontology. For this reason, if the terms are composite words, $termSim$ tokenizes the word and then returns the average similarity among the terms.

As final remark, we want to say that URBE is built on top of an UDDI implementation only for historical reasons, but the similarity engine has wider applicability as we will see in Section 4. Such a module can also be used as a stand-alone component or be embedded in complex frameworks.

3.3 USQL

USQL (Universal Service Query Language, [11]) is an XML language to express service requirements in a technology agnostic way. The language allows users to abstract the particular protocol and details used by the registry (i.e., *how*), and focus their attention on *what* services are supposed to offer. USQL is thus a language for searching services understood by different registry vendors, like

⁷ <http://wordnet.princeton.edu/>

SQL in the database world. The simplicity, expressiveness, and extensibility of the language make USQL a good solution both for experts and unskilled users. For example, users without technical skills can search for services provided by certain organizations, while more skilled users can search for services that offer particular operations. A special-purpose engine translates both the queries from users into the format imposed by the particular registry, and the responses from registry-dependent descriptions into a *generic service model* (GeSMO). This model adopts a layered structure: the lowest layer contains the concepts common to different services. Higher layers describe properties specific to particular services. This way, we have an extensible model able to capture different service types (e.g., Web services, Grid services, and P2P services) using orthogonal metrics, like semantics, QoS, trust and security, and management.

USQL queries can exploit *syntactic* information about the services, for example, their names or the names of the organizations behind them. They can also embed *semantic* data that belong to users' domain knowledge, and *QoS* elements to predicate on the non-functional requirements that the service is supposed to comply with. Obviously, we can easily mix these data to conceive complex and sophisticated queries to retrieve the services of interest.

USQL is based on a simple XML dialect to describe both required services and their QoS properties. In particular, there are elements to select services with a particular name, with a particular service description, or provided by a particular service provider. As for semantics, USQL supports different taxonomy schemes such as the *North American Industry Classification System* or the *United Nations Standard Products and Services Code System*. The user is able to specify requirements on the operations the service should provide. USQL also accepts constraints on the desired quality of service in terms of price, availability, reliability, processing time, and security. These orthogonal aspects fully support the user to retrieve services with the required functional and non-functional properties.

For example, if we want a service to send SMS messages, we might think of different properties. We can specify that interesting services must contain SMS in their name, and thus we exploit the syntactic information. Moreover, we could also exploit their semantic characterization to only discover services provided by phone companies, or require that the WSDL interface of the service we want must have a `send` method that accepts a phone number and a short message as inputs. Finally we can also say that we are only interested in cheap services by setting a maximum price.

4 Proposed solution

A set of isolated registries would force interested providers to publish their services on all of them to effectively advertise their services and increase user awareness. For this reason, we propose a flexible infrastructure that takes advantage of DIRE, URBE, and USQL to simply the way services are published over a set of registries and thus ease their retrieval.

Once a new Web service becomes available, this information is not only stored in the registry used by the provider to publish the new service, but it also forwarded to all the other registries interested in the same kind of services. This way, the provider can reduce the set of target registries to ideally a single one. In turn, even service retrieval becomes more effective: we move from a scenario where requesters have to browse different registries to find what they want, to a scenario where requesters only express their needs once, their requirements are spread around, and the information about interesting Web services is automatically moved onto their registries.

The result of this combination is shown in Figure 3. The core of the infrastructure is similar to the one adopted in DIRE, where a *communication bus*⁸ connects all the companies that own a registry. Generally speaking, every registry can be used to both publish new Web services and retrieve interesting ones. Each registry is connected to the bus by means of the *delivery manager*, which is in charge of the different registry technologies and also manages the information flow in the two directions: it communicates with the *dispatcher*, which is part of the *communication bus* and implements the mechanisms for receiving new subscriptions and notifying the publication of new Web services. The first significant addition of this paper is that the *communication bus* also relies on an extended version of URBE's *similarity engine* for the comparisons between requests (subscriptions) and available services.

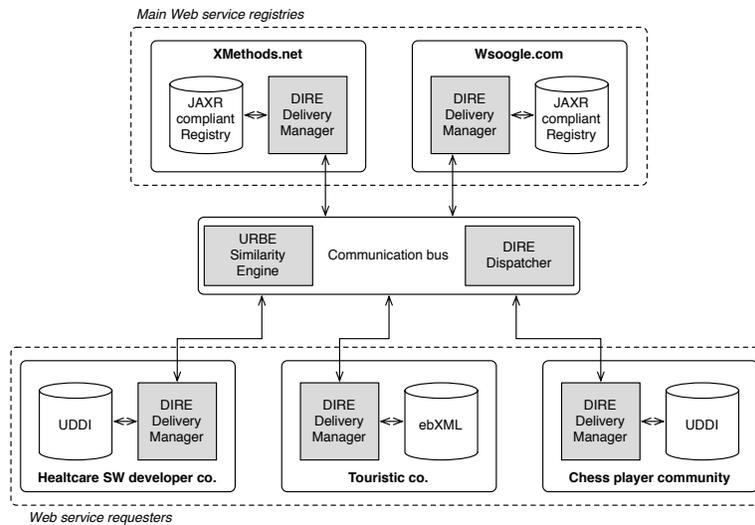


Fig. 3. Infrastructure supporting running example

⁸ We can easily assume secure and reliable interactions since the P/S communication infrastructure is in charge of it.

The figure also shows how the proposed solution works with our running example. For the sake of simplicity, we assume that XMethods and Wsoogle are connected to the communication bus via a modified version of the *delivery manager* since we are considering general purpose registries, which might host a high number of services, and the P/S infrastructure might become the bottleneck of the entire system. To avoid this, we assume the existence of different “thematic” buses (e.g., about games, health, and so on). Our general purpose registries, instead of publishing all their services onto a single bus, only notify new services on related buses. Said this, the explanation of the approach can easily consider a single bus without losing any significant detail. In addition, we also assume that all the three actors, interested in new Web services, have a proprietary service registry, along with a delivery manager properly connected to the communication bus.

The introduction of USQL and URBE aims at (i) affecting the way subscriptions can be expressed and (ii) improving the effectiveness of the filtering performed by DIRE’s *dispatcher* when it has to decide whether to forward the information about a new service or not. The next sections illustrate these two aspects in detail.

4.1 USQL-based subscriptions

This section explains how users can interact with the *delivery manager* using USQL. As stated before, the main benefits are the independence of any particular technology and the openness towards non technical users. Our additional goal is to leverage these features inside our distributed environment, which means translating USQL queries into appropriate subscriptions to support the two-phase lookup. Notice that USQL queries are nothing but verbose XML documents and space constraints force us to only show an example.

Moving back to the scenario of Section 2, chess players had discovered an interesting set of services provided by an organization called AcmeChess. Now, the community is looking for other services, and given the good reputation they would like to understand the new services provided by AcmeChess. For example, if we think of a simple facet with a tag `serviceprovider`, the filter (i.e., the XPath expression) could be: `//serviceprovider = "AcmeChess"`.

Another example considers the tourist operator that uses USQL in a smarter way and exploits the semantic facets used to describe services. If we assume the existence of a standard facet that represents shared taxonomy about travels, we can easily select all the services related to it. The subscription behind this query would predicate on the relationship between the standard Facet `travel` and the various services to allow the delivery manager to retrieve all the services in the domain of interest. Otherwise, if the ontology is more dynamic and lightweight, it can be embedded into a specific facet used to describe a particular service and the travel agency uses: `contains(//service/ontology, "traveling")`.

The third actor, that is, the health-care software development company, selects services by analyzing the interface they offer. For this reason, the first query they create analyzes the operations exposed by the different services to select the

ones relevant for their goals. For example, the USQL query in Figure 4 wants a service that gives the patient's name knowing his/her social security number. It also checks that the service only requires an integer tagged with the taxonomy's node `ssn`, and returns a string, tagged with node `surname`. This query is transformed into a filter with three parts. The first part is an XPath that analyses the WSDL facet to check whether there is an operation with an integer input and a string result. The second part checks whether the input parameter of the operation refers to taxonomy's node `ssn`, while the last part checks the result of the operation, and controls that it represents a `surname`.

```

<USQL version="1.0" xmlns="urn:sodium:USQL">
  <USQLRequest>
    <Where>
      <Service serviceType="WebService P2PService GridService">
        <Operation minDegreeOfMatch="0.75">
          <Inputs>
            <input>
              <type valueIs="contain">integer</type>
              <semantics ontologyURI="http://sodium/ontologies/healthcare">
                http://sodium/ontologies/healthcare#ssn
              </semantics>
            </input>
          </Inputs>
          <Outputs>
            <output>
              <type valueIs="contain">string</type>
              <semantics ontologyURI="http://sodium/ontologies/healthcare">
                http://sodium/ontologies/healthcare#surname
              </semantics>
            </output>
          </Outputs>
        </Operation>
      </Service>
    </Where>
  </USQLRequest>
</USQL>

```

Fig. 4. USQL query defined by the Healthcare SW company.

When a service satisfies the three parts of the filter, its description is copied onto the user's registry. If the query is issued at run-time, we assume that the user wants exact results (i.e., services that can replace a previous service without any human intervention), and the infrastructure uses the standard XPath matching technique. If the query is performed at design-time, the user may want approximate results, USQL allows us to specify the degree of matching, and URBE helps the communication bus to select services this way.

The health-care development company can also decide to include QoS requirements in its queries. For example, they can decide to only bind to services with high availability and low processing time. These concepts are supported by USQL, which also allows us to predicate on availability, response time, security level, and required price. All these elements can easily be translated into both XPath queries directly, for full matches, and complete required facets, which are then passed to URBE for the evaluation of partial matches.

4.2 Similarity-based subscriptions

The extended version of URBE's *semantic matching* provides two main features: *termSim*, to evaluate the similarity between two terms, and *facetSim*, to evaluate the similarity between two facets, which is an extension and generalization of the original match making that was limited to WSDL or SAWSDL descriptions. The infrastructure we propose exploits these two functions whenever users want to move from *exact match* comparisons to *relaxed match* ones, that is, users are satisfied, but their requirements are not totally fulfilled.

To notify the publication of a new Web service, the dispatcher was used to verify that the new description and the subscription had a perfect match. For example, if the chess player community submits a subscription with an XPath expression as `//service/type="chessgame"`, their registry will never receive Web services with a facet whose field `type` is `chess`. Since we cannot force all the actors to use the same terms, we can take advantage of function *termSim*. We introduce the clause `relaxed[sim]`, and append it to the XPath expression included in the subscription, where $sim \in [0..1]$ is a threshold that specifies the minimum admissible similarity. In the example, the subscription could be `//service/type="chessgame" relaxed[0.5]` to make the dispatcher notify the publication of new Web services with $termSim('chessgame', 'chess') \geq 0.5$. Notice that the definition of this similarity threshold is not easy for unskilled users as the average chess player. For this reason, we assume that the `relaxed` clauses can actually be set by transforming a qualitative scale (e.g., high, medium, and low) into the corresponding threshold values.

Function *facetSim* can be exploited in case the requester is skilled enough to know what a facet is (that is, the structure of an XML document used to describe a service). Thus, the requester wants a Web service that is not only related to a given type, but it is described in a very particular and technical way. This scenario is very common in autonomic systems where we need to ascribe applications with self-* properties [19]. To make the substitution possible, the substitute Web service has to expose a facet that is equal to or at least similar to the facet of the failed service. Since a WSDL description is nothing but a particular facet, a subscription of the healthcare software company could be `//service/wsd1='http://www.hcc.org/x-rayPrinter' relaxed[0.8]` where the WSDL corresponds to a Web service able to print and deliver X-rays to patients. When a company develops a new service of this kind and publishes it onto one of the general purpose registries, the dispatcher compares its WSDL with the WSDL at `http://www.hcc.org/x-rayPrinter`. If *facetSim* returns a

value greater than 0.8, the Web service is also published onto the private registry owned by the software company.

5 Related work

Our proposal can easily be compared with two wide classes of approaches: those that concentrate on service publication and discovery and those that deal with term similarity.

As for the first group, Garofalakis et al. in [1] introduce an overview of current Web service publication and discovery mechanisms and also propose a categorization. Registry technologies support the cooperation among registries, but they imply that all registries comply with a single standard and the cooperation needs a set up phase to manually define the information contributed by each registry. For example, UDDI v.3 [20] extends the replication and distribution mechanisms offered by the previous versions to support complex and hierarchical topologies of registries. It also identifies services by means of a unique key over different registries. The standard only says that different registries can interoperate, but the actual interaction policies must be defined by the developers.

Similarly, ebXML [4] is a family of standards based on XML to provide an infrastructure to ease the online exchange of commercial information. ebXML fosters the cooperation among them by means of the idea that groups of registries share the same commercial interests or are located in the same domain. One of such groups can then be seen as a single logical entity where all the elements are replicated on the different registries

METEOR-S [8] and PYRAMID-S [9] fall in the family of semantic-aware approaches for the creation of scalable peer-to-peer infrastructures for the publication and discovery of services. The semantic infrastructure allows for the implementation of different algorithms for the publication and discovery of services, but it also forbids the complete control over the registries. The semantic layer imposes too heavy constraints on publication policies and also on the way federations can evolve dynamically. METEOR-S only supports UDDI registers, while PYRAMID-S supports both UDDI and ebXML registries. They adopt ontology-based meta-information to allow a set of registries to be federated with each registry “specialized” according to one or more categories it is associated with. This means that the publication of a new service requires the meta-information needed to categorize the service within the ontology. Services are discovered by means of semantic templates that give an abstract characterization of the service and are used to query the ontology and identify the registries that contain significant information.

Term similarity has been tackled in several different ways [18]. These algorithms usually calculate such a similarity by relying on the relationships between terms defined in a reference ontology (e.g., *is-a*, *part-of*, *attribute-of*). In contrast, we compute similarity between terms according to the approach proposed by Seco et al. [21], where the authors adapt existing approaches with the assump-

tion that concepts with many hyponyms convey less information than concepts that have less hyponyms or any at all (i.e, they are leaves in the ontology).

About the similarity between whole signatures, our approach is closely related to the approaches studied for the retrieval of reusable components [22]. In this field, as stated by Zaremski and Wing, there are two types of methods to address this problem: signature matching [23] and specification matching [24]. In particular, signature matching considers two levels of similarity and introduces the *exact* and *relaxed* matching of signatures. As for services, Stroulia and Wang [25] propose an approach that also considers the description field usually included in WSDL specifications.

6 Conclusions and future work

The paper presents an innovative infrastructure for the distributed publication of Web services and for their easy retrieval. The proposal leverages previous experiences of the authors, namely DIRE and URBE, and also other initiatives (USQL) to provide a holistic solution able to govern the replication of service information by means of user requests and preferences, and also able to provide users with partial, but acceptable, solutions whose fitness is defined through semantic match making techniques. The overall framework provides users with a wide set of options.

The integrated infrastructure exists as a very first prototype, but more stable solutions are needed for its deployment in realistic settings and also for a thorough empirical evaluation of the approach. Both these directions will govern our future work. The plan is to keep working on a fully functional prototype implementation and design a complete empirical evaluation of the proposal by exploiting a fully distributed set of registries and the usual collections of public Web services as benchmarks.

Acknowledgment

This work has been supported by the following projects: Tekne (Italian FIRB), Discorso (Italian FAR), SeCSE (EC IP), and ArtDecò (Italian FIRB).

References

1. Garofalakis, J., Panagis, Y., Sakkopoulos, E., Tsakalidis, A.: Contemporary Web service discovery mechanisms. *Journal of Web Engineering* **5**(3) (2006) 265–290
2. Andrews, T., Curbera, F., Dholakia, H., Golland, Y., Klein, J., Leymann, F., Liu, K., Roller, D., Smith, D., Thatte, S., Trickovic, I., Weerawarana, S., Thatte, S.: Business Process Execution Language for Web Services Version 1.1. Specification, BEA Systems, Int'l Business Machines Corporation, Microsoft Corporation, SAP AG, Siebel Systems (2003)
3. : The UDDI Web site. (<http://uddi.xml.org>)

4. : ebXML: Electronic Business using eXtensible Markup Language. (<http://www.ebxml.org/>)
5. Martin (ed.), D.: OWL-S: Semantic Markup for Web Services. W3C Submission. <http://www.w3.org/Submission/2004/SUBM-OWL-S-20041122/> (2004)
6. WSMO Working Group: Web Service Modeling Ontology. (<http://www.wsmo.org>)
7. Farrel, J., Lausen, H.: Semantic annotations for WSDL and XML schema. <http://www.w3.org/TR/sawSDL/> (2007)
8. Verma, K., Sivashanmugam, K., Sheth, A., Patil, A., Oundhakar, S., Miller, J.: METEOR-S WSDI: A scalable p2p infrastructure of registries for semantic publication and discovery of web services. In: Information Technology and Management. Volume 6. (Jan 2005) 17 – 39
9. Pilioura, T., Kapos, G., Tsalgatidou, A.: PYRAMID-S: A scalable infrastructure for semantic web services publication and discovery. In: RIDE-DGS 2004 14th Int'l Workshop on Research Issues on Data Engineering, in conjunction with the IEEE Conf. on Data Engineering (ICDE 2004). (March 2004)
10. Eugster, P.T., Felber, P.A., Guerraoui, R., Kermarrec, A.M.: The many faces of publish / subscribe. *ACM Comput. Surveys* **35(2)** (2003) 114131
11. Tsalgatidou, A., Pantazoglou, M., Athanasopoulos, G.: Specification of the Unified Service Query Language (USQL). Technical report, (June 2006)
12. Baresi, L., Miraz, M.: A distributed approach for the federation of heterogeneous registries (2006)
13. Plebani, P., Pernici, B.: Web service retrieval based on signatures and annotations. Technical Report 2007.47, Dipartimento di Elettronica ed Informazione - Politecnico di Milano (2007)
14. Najmi (ed.), F.: Java api for xml registries (jaxr). proposed final draft. <http://java.sun.com/webservices/jaxr/> (2002)
15. Sawyer, P.: Specification language definition. Technical Report A1.D2.3, EC SeCSE Project (2006)
16. Cugola, G., Picco, G.: Reds: a reconfigurable dispatching system (2006)
17. Bianchini, D., De Antonellis, V., Pernici, B., Plebani, P.: Ontology-based methodology for e-service discovery. *Information Systems* **31(4-5)** (2006) 361–380
18. Pedersen, T., Patwardhan, S., Michelizzi, J.: WordNet::Similarity - measuring the relatedness of concepts. In: Proc. National Conf. on Artificial Intelligence, July 25-29, San Jose, California, USA. (2004) 1024–1025
19. Kephart, J.O., Chess, D.M.: The vision of autonomic computing. *IEEE Computer* **36(1)** (2003) 41–50
20. Clement, L., Hately, A., von Riegen, C., (eds.), T.R.: Universal Description, Discovery and Integration version 3.0.2. http://uddi.org/pubs/uddi_v3.htm (2004)
21. Seco, N., Veale, T., Hayes, J.: An intrinsic information content metric for semantic similarity in Wordnet. In: Proc. European Conf. on Artificial Intelligence (ECAI'04), Valencia, Spain, August 22-27, IOS Press (2004) 1089–1090
22. Damiani, E., Fugini, M.G., Bellettini, C.: A hierarchy-aware approach to faceted classification of object-oriented components. *ACM Trans. Softw. Eng. Methodol.* **8(3)** (1999) 215–262
23. Zaremski, A., Wing, J.: Signature matching: a tool for using software libraries. *ACM Trans. Softw. Eng. Methodol.* **4(2)** (1995) 146–170
24. Zaremski, A., Wing, J.: Specification matching of software components. *ACM Trans. Softw. Eng. Methodol.* **6(4)** (1997) 333–369
25. Stroulia, E., Wang, Y.: Structural and semantic matching for assessing Web-service similarity. *Int'l J. Cooperative Inf. Syst.* **14(4)** (2005) 407–438