

Discovering Services: Towards High-Precision Service Retrieval*

Abraham Bernstein¹ and Mark Klein²

¹ New York University, New York, NY 10012, USA
avi@acm.org

² MIT – CCS, 77 Massachusetts Ave, Room NE20-336, Cambridge, MA 02139, USA
m_klein@mit.edu

Abstract. The ability to rapidly locate useful on-line services (e.g. software applications, software components), as opposed to simply useful documents, is becoming increasingly critical in many domains. Current service retrieval technology is, however, notoriously prone to low precision. This paper describes a novel service retrieval approach based on the sophisticated use of process ontologies. Our preliminary evaluations suggest that this approach offers qualitatively higher retrieval precision than existing (keyword and table-based) approaches without sacrificing recall and computational tractability/scalability.

1 The Challenge: High Precision Service Retrieval

Increasingly, on-line repositories such as the World Wide Web are being called upon to provide access not just to documents that collect useful *information*, but also to *services* that describe or even provide useful *behavior*. Potential examples of such services abound:

- *Software applications* such as web services that can be invoked remotely by people or software. E.g., www.salcentral.com.
- *Software components* that can be downloaded for use when creating a new application. E.g., www.mibsoftware.com, www.compoze.com.
- *Best practice repositories* that describe how to achieve some goal. E.g., process.mit.edu/eph/, www.bmpcoe.com.
- *Individuals or organizations* who can perform particular functions, E.g., guru.com, elance.com, and freeagent.com.

As the number of such services increase it will become increasingly important to provide tools to quickly find the services they need, while minimizing the burden for those who wish to list their services with these search engines [1]. Current service retrieval approaches have, however, serious limitations with respect to meeting these

* To appear as: Bernstein, Abraham, and Mark Klein. 2002. "Discovering Services: Towards High Precision Service Retrieval" in *Proceedings of the CaiSE workshop on Web Services, e-Business, and the Semantic Web: Foundations, Models, Architecture, Engineering and Applications*. Toronto, Canada

challenges. They either perform relatively poorly or make unrealistic demands of those who wish to index or retrieve services. This paper first reviews these approaches and then presents as well as evaluates a novel service retrieval approach based on the sophisticated use of process ontologies. It closes with a discussion of open challenges for future work.

2 The State of the Art

Current service retrieval technology has emerged from several communities. The information retrieval community has focused on the retrieval of natural language service descriptions and has as a result emphasized keyword-based approaches. The software agents and distributed computing communities have developed simple ‘table-based’ approaches for ‘matchmaking’ between tasks and on-line services. The software engineering community has developed by far the richest set of techniques for service retrieval [2]. We can get a good idea of the relative merits of these approaches by placing them in a *precision/recall* space (Figure 1):

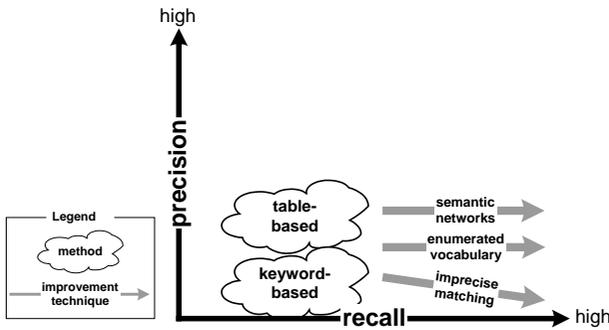


Fig. 1. State of the art in service retrieval

Most search engines, including service repositories such as www.salcentral.org, look for items that contain the *keywords* in the query, which are sometimes prioritized using techniques such as TFIDF to increase effective precision [3]. Keyword-based approaches are, however, prone to both low precision, as irrelevant items may, e. g., contain the keyword, and imperfect recall due issues such as the use of synonyms (sometimes addressed with pre-enumerated vocabularies [4], semantic nets [5] and partial matching). The key underlying problem is that keywords are a poor way to capture the semantics of a query or item.

Table-based approaches [6] [7] [8] [9] [10] [11] typically use a fixed number of attribute value pairs describing the properties of a service. Figure 2, e.g., shows a table-based model for an integer averaging service:

Description	find the average of a list of integers
Input	Integers
Output	Real
Duration	number of inputs * 0.1 msec

Fig. 2. A table-based description of an integer sorting service

Both items and queries are described as tables: matches represent items whose property values match those in the query. All the commercial service search technologies we are aware of (e.g., Jini™, eSpeak, Salutation, UDDI/WSDL, [12]) use this approach. The more sophisticated search tools emerging from the research community [13] [14] use ontologies and semantic nets to increase recall, e.g. returning a match if the input type of a service is equal to *or* a generalization of the input type specified in the query. Table-based models, however, do little to increase precision because of the impoverished range of information they capture, as they typically include a detailed description of how to *invoke* the service (i.e., parameter types, return types, etc.), but don't describe what the service actually *does*. This information is of limited value for search purposes because services with different goals (e.g., services that compute averages, medians, quartiles, etc.) can share similar call signatures.

Other approaches (such as deductive retrieval [15] or execution-based retrieval [16]) are usually only suitable for limited application domains, as they are, typically, too complex (both from a computational and a usability perspective).

3 Our Approach: Exploiting Process Ontologies

Our challenge can thus be framed as being able to capture enough service and query semantics to substantively increase precision without reducing recall or making it unrealistically difficult for people to express these semantics. *Our central claim is that these goals can be achieved through the sophisticated use of process ontologies* [17]. In our approach, the salient behavior of a service is captured using process models, and these process models, as well as their components (subtasks, resources, etc.), are placed in the appropriate locations in the process ontology. Queries can then be defined (using a *process query language* – PQL) to find all the services whose process models include a given set of entities and relationships. The greater expressiveness of process models, as compared to keywords or tables, offers the potential for substantively increased retrieval precision, at the cost of requiring that services be modeled in this more formal way. As we will see below, our preliminary evaluations suggest that the process-based approach offers qualitatively increased retrieval precision, and we will argue that this can be achieved with a reasonable expenditure of service modeling effort. Our approach has the functional architecture shown in Figure 3, which we will consider below.

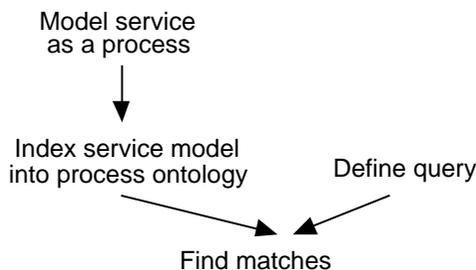


Fig. 3. Functional architecture

3.1 Modeling Services as Process Models

The first step in our approach is to capture service behavior as process models. Why process models? To understand this choice, we need to understand more precisely the causes of imperfect precision. One cause is that a component of the service model is taken to have an unintended *role*. For example, a keyword-based query to find mortgage services that deal with “payment defaults” (a kind of exception) would also match descriptions like “the payment defaults to \$100/month” (an attribute value). The other cause for false positives occurs when a service model is taken to include an unintended *relationship* between components. For example, we may be looking for a mortgage service where insurance is provided for payment defaults, but a keyword search would not distinguish this from a service that provides insurance for the home itself. The trick to increasing retrieval precision, therefore, comes down to ensuring that important roles and relationships are made explicit in both the query and the service model, so unintended meanings (and therefore false positives) can be avoided. Process modeling languages are well suited for this as they have been designed to capture the essence of different behaviors in a compact intuitive way, and have become ubiquitous for a very wide range of uses.

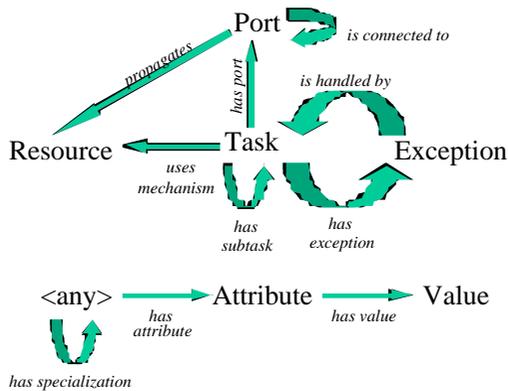


Fig. 4. Process model formalism

We use for our purposes a process modeling formalism (see Figure 4) that, similar to other processes modeling languages, includes the following components:

- *Attributes*: capture such information as a textual description, typical performance values (e.g., execution time), etc.
- *Decomposition*: A process can be modeled as a collection of processes that can in turn be broken down (“decomposed”) into sub-processes.
- *Resource Flows*: All process steps can have input and output *ports* through which *resources* flow allowing us to model consumed, and produced resources.
- *Mechanisms*: Processes can be annotated with the resources they *use* (as opposed to consume or produce). E.g., the Internet can serve as a mechanism for a process.

- *Exceptions*: Processes typically have characteristic ways they can fail and associated schemes for anticipating and avoiding or detecting and resolving them. Our approach captures these schemes by annotating processes with their characteristic ‘exceptions’, and mapping these exceptions to processes describing how these exceptions can be handled [18].

Let us consider a simple example to help make this more concrete (Figure 5):

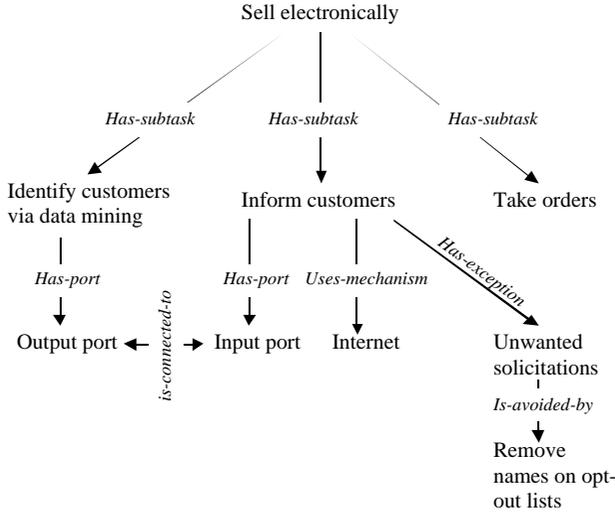


Fig. 5. An example of a process-based service model

This represents the process model for a service for selling items electronically. The plain text items represent entities (such as exceptions, ports, etc.), while the italicized items represent relationships between these entities. The substeps in this service model include ‘identify potential customers via data mining’, ‘inform customers’ (which uses the Internet as a mechanism), and ‘take orders’. The potential exception of sending out unwanted solicitations is avoided by filtering out the names of individuals who have placed their names on ‘opt-out’ lists. Each of the entities can have attributes (not shown) that include their name, description, and so on.

Formally, any database of process descriptions (using the formalism above) can be defined as a typed graph:

$$Ont (Entities, Relationships) \tag{1}$$

where *entities* are the nodes in the graph and *relationships* are the graph edges. Furthermore, the following specifications apply:

- A node can only have one type (\oplus denoting exclusive-or):

$$x \in Entities \equiv (x \in Task) \oplus (x \in Resource) \oplus (x \in Port) \oplus (x \in Exception) \oplus (x \in Attribute) \oplus (x \in Value) \tag{2}$$

- A relationship can only have one type and it connects nodes of certain types:

$$\begin{aligned}
r(x, y) \in \text{Relationships} \equiv & \\
& [(r = \text{Has_specialiation}) \wedge (x \in \text{Entities}) \wedge (y \in \text{Entities}) \wedge ((x, y) \in \text{Has_specialization})] \oplus \\
& [(r = \text{Has_subtask}) \wedge (x \in \text{Task}) \wedge (y \in \text{Task}) \wedge ((x, y) \in \text{Has_subtask})] \oplus \\
& [(r = \text{Has_port}) \wedge (x \in \text{Task}) \wedge (y \in \text{Port}) \wedge ((x, y) \in \text{Has_port})] \oplus \\
& [(r = \text{Uses_mechanism}) \wedge (x \in \text{Task}) \wedge (y \in \text{Resource}) \wedge ((x, y) \in \text{Uses_mechanism})] \oplus \\
& [(r = \text{Propagates}) \wedge (x \in \text{Port}) \wedge (y \in \text{Resource}) \wedge ((x, y) \in \text{Propagates})] \oplus \\
& [(r = \text{Is_handled_by}) \wedge (x \in \text{Exception}) \wedge (y \in \text{Task}) \wedge ((x, y) \in \text{Is_handled_by})] \oplus \\
& [(r = \text{Has_exception}) \wedge (x \in \text{Task}) \wedge (y \in \text{Exception}) \wedge ((x, y) \in \text{Has_exception})] \oplus \\
& [(r = \text{Has_subtask}) \wedge (x \in \text{Task}) \wedge (y \in \text{Task}) \wedge ((x, y) \in \text{Has_subtask})] \oplus \\
& [(r = \text{Has_attribute}) \wedge (x \in \text{Entities}) \wedge (y \in \text{Attribute}) \wedge ((x, y) \in \text{Has_attribute})] \oplus \\
& [(r = \text{Has_value}) \wedge (x \in \text{Attribute}) \wedge (y \in \text{Value}) \wedge ((x, y) \in \text{Has_value})] \oplus
\end{aligned} \tag{3}$$

This representation is similar, and equivalent in expressiveness, to other full-fledged process modeling languages (e.g. IDEF [19], PIF [20], PSL [21] and CIMOSA [22]) and substantially more expressive than the keyword and table-based languages used in previous service retrieval efforts, by virtue of adding the important concepts of resource flows, task decompositions, and exceptions. It does not however, include primitives oriented at expressing control semantics, i.e. that describe *when* each subtask gets enacted. Such primitives were excluded for two reasons. First, most of the variation between process modeling languages occurs when representing control semantics, and we wanted a formalism to which a wide range of existing process models could easily be translated. Second, our experience to date that most service queries are concerned with *what* a process does, rather than *when* the parts of the process gets enacted.

Modeling service behaviors as process models of course involves manual effort, but we argue this need not be a major barrier. Because process formalisms are so widely used, many services already have process models defined for them, e.g., as part of their specification. The expertise needed to create such models is widely available. Process ontologies (see below) can reduce the modeling effort involved. Also, service providers will likely be motivated to create such process models, since they often differentiate themselves in the marketplace by how they provide their services, and process models make this explicit. Process models, finally, enable important uses other than search, such as automatic service composition.

3.2 Indexing Service Models into the Process Ontology

The second step of our approach is to index service models into a process ontology in order to facilitate later retrieval. An ontology consists, in general, of a hierarchy of entity descriptions ranging from the abstract at one end to the specific at the other. Items with similar semantics (e.g. processes with similar functions) appear close to each other, the way books with similar subjects appear close to each other in a library. Indexing a service comes down to placing the associated process model, as well as all of its components (attributes, ports, dependencies, subtasks and exceptions) on the appropriate branch in the ontology. Using an ontology is valuable for several reasons. It can reduce the burden of modeling a service, since one need only find the most similar process model in the ontology and then modify it to model a new service. Ontologies can increase recall, since similar services are co-located, one is apt to find

relevant services simply by browsing the ontology near the matches one has already found. In addition, an ontology helps us find matches that are described using different, but semantically equivalent, terminology.

We build for this purpose on the MIT Process Handbook project, a process ontology, which has been under development at the Center for Coordination Science (CCS) for the past ten years [23] [24]. The growing Handbook database currently includes roughly 5000 process descriptions ranging over such areas as supply chain logistics, hiring, etc. The Handbook project has developed sophisticated tools that allow a knowledgeable user to index a process model in a matter of minutes. We believe that the Handbook ontology represents an excellent starting point for indexing many services because it is focused on business processes, which is what a high proportion of such services are likely to address.

3.3 Defining Queries

It is of course imaginable that we could do without queries entirely once services have been indexed into an ontology. One could simply browse the ontology to find the services that one is interested in, as in [25]. Our experience suggests however that browsing can be slow and difficult for all except the most experienced users. This problem is likely to be exacerbated when, as with online services, the space of services is large and dynamic. To address this challenge we have defined a query language called PQL (the Process Query Language) designed for retrieving process models indexed in an ontology. Process models can be straightforwardly viewed as entity-relationship diagrams made up of entities like tasks characterized by attributes and connected by relationships like ‘has-subtask’. PQL queries are built up as combinations of three primitive clause types that check for these elements:

- Entity <entity> isa <entity type>
- Relation <source entity> <relationship type> <target entity> [*]
- Attribute <attribute> of <entity> {includes | equals} <value>

The ‘entity’ clause matches any entity of a given type (the entity types include task, resource, port and so on). The ‘relation’ clause matches any relationship of a given type between two entities (the relationship types include has-subtask, has-specialization, has-port, and so on). The optional asterisk finds the transitive closure of this relationship. The ‘attribute’ clause looks for entities with attributes that have given values. Any bracketed item <> can be replaced by a variable (with the format ?<string>) that is bound to the matching entity and passed to subsequent query clauses.

The ‘When’-clause allows to group clauses into sub-queries:

- When {exists | does-not-exist} <query>

Let us consider a simple example to help make this more concrete. The query below searches for a sales service that uses the internet to inform customers:

```
attribute "Name" of ?sell includes "sell"
relation ?sell has-specialization ?process *
when exists (relation ?process has-subtask ?subtask *
  attribute "Name" of ?subtask includes "inform"
  attribute "Description" of ?subtask includes "internet")
```

The first clause searches for a processes in the ontology whose name includes “sell”, the second finds all specializations of this, and the third checks if any subtasks of these services are “inform” processes with “internet” in their description. A PQL query is thus equivalent to a typed sub-graph pattern, and any search for a process model can then be treated as finding the nodes of type *task*, which match the graph pattern that represents the query.

The three clause types of PQL, and their variants, can be formalized as follows:

Relation x rel-type y

is defined as: $(x,y) \in \text{rel_type}$, where $\text{rel_type} \in \{\text{has_specialization, has_subtask, ...}\}$

Relation x rel-type y *

is formalized using a fixpoint/recursive definition, as:

$((x, y) \in \text{rel_type}) \vee (\exists z : ((x,z) \in \text{rel_type}) \wedge \text{rel_type}(z,y)^*)$

Entity entity isa entity-type

is defined as:

$\text{entity} \in \text{entity_type}$, where $\text{entity_type} \in \{\text{Task, Port, Resource, Exception, Attribute, ...}\}$

Attribute attribute of entity equals value

is shorthand for two relationships, as follows:

$\text{has_attribute}(\text{entity, attribute}) \wedge \text{has_value}(\text{attribute, value})$

Attribute attribute of entity includes value

is defined as:

$\text{has_attribute}(\text{entity, attribute}) \wedge \text{has_value}(\text{attribute, v1}) \wedge \text{IsSubString}(\text{value, v1})$

Note that PQL includes built-in functions, such as *IsSubString*, comparable to those in other query languages such as SQL. If any of the parameters to a predicate are preceded by a question mark (e.g., ?y), then it denotes a variable that needs to be bound to a value from the database/model that can fulfill its place.

The ‘when’ construct serves two roles. If used with the “exists” operator then it simply groups sub-queries in an intuitive way, and does not add any expressive power to PQL. For example:

Relation ?x Has_subtask ?y

When exists ((Relation ?y Has_subtask a))

is formalized equivalently with or without the “when” operator, as:

$\text{Has_subtask}(?x, ?y) \wedge \text{Has_subtask}(?y, a)$

If the when-statement is used with the “does-not-exist” option then it will only return a result if <query> does not. This introduces a form of *negation* into PQL, so:

When does-not-exist query

is defined as $\neg \exists x_{i=1..k} : x_{i=1..k} \in \text{Entities} : \langle \text{query} \rangle$

where: $x_{i=1..k}$ are the unbound variables in query.

The question of how to add negation to a query language is a non-trivial issue, as it may have major implications on its computational tractability. As will become obvious in section 5 below, the type of negation introduced here is consistent with an inflationary fixpoint approach, ensuring that the resulting language is bounded by polynomial time.

As a final example, let us consider how our original example PQL query is formalized:

```

Has_attribute(?sell, "Name") ∧ Has_value("Name", ?v1) ∧ IsSubString("sell", ?v1) ∧
Has_Specialization(?sell, ?process) ∧
Has_attribute(?process, "service?") ∧ Has_value("service?", "yes") ∧
( Has_subtask(?process, ?subtask) ∧
  Has_attribute(?subtask, "Name") ∧ Has_value("Name", ?v2) ∧ IsSubString("inform", ?v2) ∧
  Has_attribute(?subtask, "Description") ∧ Has_value("Description", ?v3) ∧
  IsSubString("inform", ?v3) )

```

PQL has been used successfully to represent a wide range of queries drawn from many different domains. Some other examples include “find a loan process that uses the internet, takes real estate as collateral, and has loan default insurance”, “find all processes that take oil as an input and are prone to cause environmental damage”, and so on. Our preliminary assessment is that PQL is sufficiently expressive to capture all queries describable in process-oriented terms.

3.4 Finding Matches

The algorithm for retrieving matches given a PQL query is straightforward. The clauses in the PQL query are tried in order, each clause executed in the variable binding environment accumulated from the previous clauses. The bindings that survive to the end represent the matching services. While we have not yet evaluated PQL’s performance in detail yet, we do show (see below) that queries are within polynomial time complexity.

4 Empirical Evaluation

An initial version of the PQL interpreter has been implemented, and we have performed some preliminary evaluations of its precision and recall compared to existing (keyword and table-based) approaches. The following scheme was used for all of the evaluations described below. The roughly 5000 processes in the Process Handbook process ontology were treated as service models, which is reasonable since they all represent functions used in business contexts and many could imaginably be performed remotely. We then defined keyword and process-based queries that use the same keywords, operate over the same database of service models, and differ only in whether they use the role and relationship information encoded in the service models. We did not define a separate set of table-based service models and queries for this evaluation because, from the standpoint of retrieval precision, the keyword and table-based approaches are equivalent. The fact that table-based models differentiate name and description attributes does not help since descriptions almost invariably reprise the keywords included in the service name, and none of the queries we used made use of I/O specifications. In any case, if we had used queries that refer to such I/O specs, it would not change the relative precision of table- and process-based queries since both can use I/O information. We tested simple keyword search as well as TFIDF, the latter because its potentially greater effective precision makes it a dominant scheme for keyword-based search. All the queries in our evaluation had perfect recall, because of the consistency in the use of keywords in the process descriptions. While we

do not anticipate that process-based search will differ significantly in recall from keyword and table-based, this remains a subject for future evaluation. The queries below, clearly, are only illustrative, since a complete evaluation would require executing a representative range of many queries.

Since our goal was to determine whether process-based retrieval improves on existing approaches, our evaluation focused on the value of the additional information captured by process-based service models as compared to keyword- and table-based models. This additional information falls into five categories: task decompositions, port connectivity, exception handling, task mechanisms, and specializations. We examine each category in the sections below.

4.1 Task Decomposition

Our process-based service model allows us to explicitly describe the subtasks that make up a service's behavior. This can help avoid confounding information that refers to different subtasks. Imagine, for example, that we are searching for a sales service that informs customers using the Internet. We can frame this query as follows:

Table 1. Query types and actual Queries

<i>Type</i>	<i>Query</i>
Keyword-based	“Sell” “inform” “internet”
Process-based	attribute "Name" of ?service includes "sell" when exists (relation ?service has-subtask ?subtask * attribute “Name” of ?subtask includes "inform" attribute ?attr of ?subtask includes “internet”)

The keyword and table-based service models are not able to distinguish cases where “inform” and “internet” (or their synonyms) belong to the same subtask from cases where these keywords belong to different subtasks (and thus are probably not relevant). We would thus predict false positives and therefore lower precision for these approaches, and this is in fact what happens. There were 13 correct matches for this query, including such processes as “Sell travel services via electronic auction”, “Sell books via electronic store” and so on. The PQL query had 13 correct matches out of the 18 it returned, for a precision of 72%. A simple keyword-based search had 280 returns, for a precision of roughly 5%. TFIDF did not improve much upon simple keyword search in this case: its precision reached a maximum of 6% (at match 163), and its overall precision was lower because it allowed partial matches and therefore generated more total returns,

4.2 Task Mechanisms

Our process-based service model allows us to describe the mechanisms used by a task, thereby avoiding false positives due to the appearance of the same keyword with a different role. We can, for example, refine the PQL query given above so that it only matches services where the keyword “internet” appears as a mechanism (the added clauses are bold type):

```

attribute "Name" of ?service includes "sell"
when exists (relation ?service has-subtask ?subtask *
  attribute "Name" of ?subtask includes "inform"
  relation ?subtask uses-mechanism ?mechanism
  attribute "Name" of ?mechanism includes "Internet")

```

This query had 13 correct matches, as above, but this time out of 16 responses, for an improved precision of 81%.

4.3 Specialization

The has-specialization relationship enabled by our inclusion of a process ontology can be used to avoid false positives by ensuring that the service, and its components, have the semantics that we desire. For example, we can use this to refine the query presented above to only accept services whose subtask is a specialization of the generic “Inform” task, thereby pruning out services with subtasks that include the string “inform” in their name for unrelated reasons (e.g. the subtask named “Collect configuration information using Internet”):

```

attribute "Name" of ?service includes "sell"
when exists (relation ?service has-subtask ?subtask *
  attribute "Name" of ?subtask includes "inform"
  relation ?subtask uses-mechanism ?mechanism
  attribute "Name" of ?mechanism includes "Internet"
  attribute "Name" of ?class equals "Inform"
  relation ?class has-specialization ?subtask)

```

With this refinement, the query returns 13 correct matches out of 13 total, for an accuracy of 100%.

Similar ontologies have of course been made available for table-based service retrieval engines. UDDI, for example, provides the UNSPSC taxonomy of product and service categories and the NAICS taxonomy of industry codes, among others. The particular value of the ontology we utilize is that it captures functions that a business might require at a much finer grain than the taxonomies mentioned above. We believe this will be helpful for service retrieval since many queries will, no doubt, be looking for services to support business functions. At the time of writing our database did not categorize services using these other taxonomies, so we were unable to evaluate their relative merits.

4.4 Exception Handling

Our process-based service model allows us to explicitly delineate the exceptions faced by a service, as well as the handlers available for dealing with each exception. Imagine, for example, that we wish to find a sales service that informs customers via the internet but avoids the exception of sending unwanted solicitations (e.g. by filtering out the names that appear on “opt-out” lists). We can, for this purpose, refine the query described above as follows shown in Table 2.

We would expect the keyword- and table-based models to incur false positives by finding services that have the same keywords in different roles, that have that exception but do not have a handler for it, or that use a different handler (e.g. allowing

recipients to remove their name from subsequent solicitations) for the same exception. In this case, there was one correct match. PQL returned only that item, for a precision of 100%. Keyword-based search returned 188 matches (0.53% precision), and TFIDF did not do any better, returning 248 documents, with a maximum precision at 0.4%.

Table 2. Query types and Query for exception handling Query

<i>Type</i>	<i>Query</i>
Keyword-based	“Sell” “inform” “internet” “avoid” “unwanted” “opt-out”
Process-based	attribute "Name" of ?service includes "sell" when exists (relation ?service has-subtask ?subtask * attribute “Name” of ?subtask includes "inform" relation ?subtask uses-mechanism ?mechanism attribute “Name” of ?mechanism includes “Internet” attribute “Name” of ?class equals “Inform” relation ?class has-specialization ?subtask relation ?subtask has-exception ?exception attribute “Name” of ?exception includes “unwanted” relation ?exception is-avoided-by ?handler relation ?attr of ?handler includes “opt-out”)

4.5 Port Connectivity

The final category of information uniquely provided by process-based service models is port connectivity, which captures the resource flow relationships between tasks. We may, for example, want a service that generates the lists of potential customers to inform by applying data-mining techniques, which implies that the output of a data mining subtask is an input to the inform customers subtask. This would imply a PQL query like the following:

```

attribute "Name" of ?sell includes "sell"
when exists (relation ?process has-subtask ?sub1
  attribute "Name" of ?sub1 includes "inform"
  relation ?sub1 has-port ?port1
  entity ?port1 isa input-port
  relation ?process has-subtask ?sub2
  attribute "Name" of ?sub includes “mining”
  relation ?sub2 has-port ?port2
  entity ?port2 isa output-port
  relation ?port1 is-connected-to ?port2)
    
```

A query like this can avoid false positives wherein a data-mining subtask exists in the service model, but it does not provide information to the inform customers step. The data-mining subtask may be applied instead, for example, to the database of sales generated by this service. We would therefore expect the keyword- and slot-based retrieval queries to demonstrate lower precision than PQL. At the time of writing we were unable to evaluate this because the Process Handbook ontology did not include sufficient port connectivity information; this will be addressed in future work.

While a wider range of queries and services needs to be evaluated, these test cases strongly suggest that the greater expressiveness of process-based service models can

in fact result in qualitatively higher retrieval precision. Even a PQL query that only took advantage of the subtask relationships in the process-based models produced retrieval precision more than 10 times greater than keyword-based approaches. The relative advantage of PQL, moreover, increased radically as the number of relationships specified in the query increased.

4.6 Retrieval Complexity

One of the major considerations for any retrieval capability is that it must return answers in a timely way. The speed at which queries get returned should be comparable to that of existing document retrieval mechanisms, which manage to search millions of documents in seconds. Even though our experience with the prototype implementation has been favorable (i.e., queries generally take several seconds at most without exploiting well-known performance-enhancing techniques such as query optimization) there is still the question of how the performance will scale with the size of the database. As we have shown elsewhere [26] PQL can be mapped to `DATALOG \neg` , which has been shown to be bounded by polynomial time [27]. Consequently, we have every reason to believe that PQL will perform comparably to other database query languages.

5 Contributions of This Work

High retrieval service precision is widely recognised as a critical enabler for important uses that range from finding useful software components or applications, to finding best practice models, to tracking down people or organisations with the skills you need. Our work can be viewed as representing a new class of service retrieval technique that helps achieve these goals (Figure 6).

Our evaluations to date suggest that process-based queries produce retrieval precision qualitatively greater than that of existing service retrieval approaches, while retaining polynomial complexity for query enactment. This work represents a significant contribution to work on pattern matching over generic graph-like representations, such as graph grammars [28], object-oriented query languages [29], and XQuery (a query language for XML; see <http://www.w3.org/TR/xquery/>). The unique value of our work comes from exploiting the more constrained semantics of process models to enable higher-precision service retrieval.

Another line of related work comes from the semantic web community. These efforts use RDF [30] (or DAML+OIL [31]) to model services (e.g., [32]), and apply standard RDF (or DAML+OIL) query engines such as RDFSuite [33] or TRIPLE [34] to retrieve services. This work differs from ours in several important respects. Our work takes advantage of a more expressive process model, which as we have seen potentially translates into higher retrieval precision. All of these approaches capture task inputs and outputs. Only DAML-S [35] and our approach, however, captures *how* a service works (i.e. the subtasks in the service process model), and only our approach captures exceptions and their handlers. Another important difference concerns the nature of the process ontology used to index service models. The Handbook ontology,

unlike that underlying RDF and DAML+OIL, allows one to "default" (i.e. delete) inherited features [36]. It has been shown that this form of inheritance is easier to understand by non-specialists [37], thereby making it easier for them to classify services. It is our hope that the insights we are exploring will lead to the enhancement of the semantic web standards and the ability to exploit these ideas using off-the-shelf query engines.

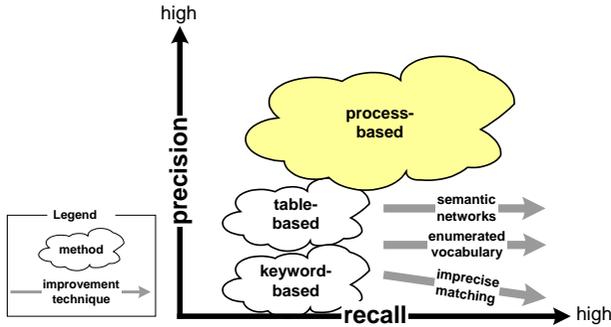


Fig. 6. The contribution of process-based service retrieval technology

6 Next Steps

While our preliminary results are promising, many important challenges remain. First and foremost, PQL needs to be evaluated for a wide range of queries in order to assess its precision and recall performance as compared to existing approaches, and to suggest refinements in the language and associated interpreter. Second, we hope to explore recall-enhancing techniques such as ontology-based query rewriting. Third, there is a need to minimize the manual effort involved in listing new services with the search engine by developing automatic indexing techniques. And finally, we will have to develop intuitive user interfaces for PQL.

References

1. J. Y. Bakos, "Reducing Buyer Search Costs: Implications for Electronic Marketplaces," *Management Science*, vol. 43, 1997.
2. H. Mili, F. Mili, and A. Mili, "Reusing software: issues and research directions," *IEEE Transactions on Software Engineering*, vol. 21, pp. 528-62, 1995.
3. G. Salton and M. J. McGill, *Introduction to modern information retrieval*. New York: McGraw-Hill, 1983.
4. R. Prieto-Diaz, "Implementing faceted classification for software reuse," *12th International Conference on Software Engineering*, vol. 9, pp. 300-4, 1990.
5. B. Magnini, "Use of a lexical knowledge base for information access systems," *International Journal of Theoretical & Applied Issues in Specialized Communication*, vol. 5, pp. 203-228, 1999.

6. S. Henninger, "Information access tools for software reuse," *Journal of Systems & Software*, vol. 30, pp. 231-47, 1995.
7. C. Fernandez-Chamizo, P. A. Gonzalez-Calero, L. Hernandez-Yanez, and A. Urech-Baque, "Case-based retrieval of software components," *Expert Systems with Applications*, vol. 9, pp. 397-421, 1995.
8. M. G. Fugini and S. Faustle, "Retrieval of reusable components in a development information system," presented at Second International Workshop on Software Reusability, 1993.
9. P. Devanbu, R. J. Brachman, P. G. Selfridge, and B. W. Ballard, "LaSSIE: a knowledge-based software information system," *Communications of the ACM*, vol. 34, pp. 34-49, 1991.
10. ESPEAK, "Hewlett Packard's Service Framework Specification,".: HP Inc, 2000.
11. R. Raman, M. Livny, and M. Solomon, "Matchmaking: an extensible framework for distributed resource management," *Cluster Computing*, vol. 2, pp. 129-38, 1999.
12. G. G. Richard, "Service advertisement and discovery: enabling universal device cooperation," *IEEE Internet Computing*, vol. 4, pp. 18-26, 2000.
13. K. Sycara, M. Klusch, S. Widoff, and J. Lu, "Matchmaking Among Heterogeneous Agents on the Internet," presented at AAAI Symposium on Intelligent Agents in Cyberspace, 1999.
14. D. Fensel, "An Ontology-based Broker: Making Problem-Solving Method Reuse Work," presented at Workshop on Problem-Solving Methods for Knowledge-based Systems at the 15th International Joint Conference on AI (IJCAI-97), Nagoya, Japan, 1997.
15. S. Meggendorfer and P. Manhart, "A Knowledge And Deduction Based Software Retrieval Tool," presented at 6th Annual Knowledge-Based Software Engineering Conference, 1991.
16. A. Podgurski and L. Pierce, "Retrieving reusable software by sampling behavior," *ACM Transactions on Software Engineering & Methodology*, vol. 2, pp. 286-303, 1993.
17. M. Klein and A. Bernstein, "Searching for Services on the Semantic Web using Process Ontologies," presented at The First Semantic Web Working Symposium (SWWS-1), Stanford, CA USA, 2001.
18. M. Klein and C. Dellarocas, "A Knowledge-Based Approach to Handling Exceptions in Workflow Systems," *Journal of Computer-Supported Collaborative Work. Special Issue on Adaptive Workflow Systems.*, vol. 9, 2000.
19. NIST, "Integrated Definition for Function Modeling (IDEF0)," National Institute of Standards and Technology, Federal Information Processing Standards (FIPS) Publication 183, December 21, 1993 1993.
20. J. Lee and M. Gruninger, "The Process Interchange Format "The Process Interchange Format and Framework v.1.2," *Knowledge Engineering Review*, 1998.
21. C. Schlenoff, M. Gruninger, M. Ciocoiu, and Jintae Lee, "The essence of the process specification language," *Transactions of the Society for Computer Simulation*, vol. 16, pp. 204-16, 1999.
22. K. Kosanke, *CIMOSA: Open System Architecture for CIM*: Springer Verlag, 1993.
23. T. W. Malone and K. Crowston, "The interdisciplinary study of coordination," *ACM Computing Surveys*, vol. 26, pp. 87-119, 1994.
24. T. W. Malone, K. Crowston, J. Lee, B. Pentland, C. Dellarocas, G. Wyner, J. Quimby, C. Osborn, A. Bernstein, G. Herman, M. Klein, and E. O'Donnell, "Tools for inventing organizations: Toward a handbook of organizational processes," *Management Science*, vol. 45, pp. 425-443, 1999.
25. L. Latour and E. Johnson, "Seer: a graphical retrieval system for reusable Ada software modules," presented at Third International IEEE Conference on Ada Applications and Environments, 1988.
26. A. Bernstein and M. Klein, "Towards High-Precision Service Retrieval," presented at The International Semantic Web Conference (ISWC-02), Sardinia Italy, 2002.

27. S. Abiteboul, R. Hull, and V. Vianu, *Foundations of Databases*. Reading, MA: Addison-Wesley Publishing, 1995.
28. G. Rozenberg, "Handbook of Graph Grammars and Computing by Graph Transformation," vol. 3: World Scientific, 1999.
29. W. Kim, *Introduction to Object-Oriented Databases*. Cambridge, MA: MIT Press, 1990.
30. O. Lassila and R. R. Swic, "Resource Description Framework (RDF) Model and Syntax Specification," W3C - World Wide Web Consortium, Cambridge, MA, W3C Recommendation 22 February 1999 1999.
31. D. Connolly, F. van Harmelen, I. Horrocks, D. L. McGuinness, P. F. Patel-Schneider, and L. A. Stein, "DAML+OIL (March 2001) Reference Description," W3C - World Wide Web Consortium, Cambridge, MA, Note 18 December 2001 2001.
32. D. Trastour, C. Bartolini, and J. Gonzalez-Castillo, "A Semantic Web Approach to Service Description for Matchmaking of Services," HP Laboratories Bristol, Bristol, UK, Technical Report July 30 th , 2001 2001.
33. S. Alexaki, V. Christophides, G. Karvounarakis, D. Plexousakis, and K. Tolle, "The RDFSuite: Managing Voluminous RDF Description Bases," presented at Second International Workshop on the Semantic Web (SemWeb 2001), Hongkong, 2001.
34. M. Sintek and S. Decker, "TRIPLE - A Query, Inference, and Transformation Language for the Semantic Web," presented at International Semantic Web Conference (ISWC), Sardinia, Italia, 2002.
35. The DAML Services Coalition, A. Ankolekar, M. Burstein, J. R. Hobbs, O. Lassila, D. L. Martin, S. A. McIlraith, S. Narayanan, M. Paolucci, and T. Payne, "DAML-S: Semantic Markup For Web Services," presented at The First Semantic Web Working Symposium, Stanford University, California, USA, 2001.
36. B. Grosz and T. C. Poon, "Representing Agent Contracts with Exceptions using XML Rules, Ontologies, and Process Descriptions," MIT - Sloan School of Management, Cambridge, MA, Submitted Working Paper April 26, 2002 2002.
37. A. MacLean, K. Carter, L. Löfstrand, and T. Moran, "User-tailorable Systems: Pressing the Issues with Buttons," presented at Human Factors in Computing Systems, Seattle, Washington, 1990.