

# Grid Service Discovery in the Financial Markets Sector

---

David Bell<sup>1</sup> and Simone A. Ludwig<sup>2</sup>

<sup>1</sup>Department of Information Systems and Computing, Brunel University, UK

<sup>2</sup>School of Computer Science, Cardiff University, UK

Investment Banking requires a diverse system set in supporting a range of markets from bonds to trading options on weather. The challenge to this community is the ability to adapt to new business requirements in an effective manner, utilizing their network of capabilities in a flexible, dynamic way. A semantic approach to discovery can be used in a pragmatic, practical manner. The use of richer explicit knowledge, that is system readable, provides the basis for discovering capabilities on this exemplar Business Grid — “the grid of services”. This design research project focuses on the utilization of disparate knowledge during discovery.

*Keywords:* grid services, ontology, discovery.

## 1. Introduction

Investment Banking software is a diverse domain to investigate — with many market-based categories ranging from Mergers to the Global Markets. The product range include Foreign Exchange, Interest Rates, Fixed Income (Bonds), Commodities (from Oil to Weather) and Derivatives. The Derivative product uses one or more of the base products as an underlying; a complex example being a total return swap which combines the characteristics of a bond with an interest stream of a swap. The trade then embarks on a lifecycle involving state transitions caused by various process enactments. Processes range from pre-trade validation between organizations to trade maturation. The resulting inventory contains many systems that focus on particular products, processes and geographical segments. Technological challenges arise when new products combine with existing ones or with synergistic process deployment. The challenge to these organizations is to be able to adapt to new business requirements in an effective way, utilizing these capabilities in a flexi-

ble manner. This research aims to address these challenges by loosening the coupling to the systems capability inventory, allowing a more configurable approach to utilization.

Grid infrastructure, with its Web Services layer, provides an infrastructure for exposing stateful organizational capabilities. The logical next step is to identify practical methods for the discovery of these capabilities when supporting this ad-hoc business capability identification. Existing methods such as UDDI (Universal Description Discovery and Integration) [2] and MDS (Monitoring and Discovery Service) [1] were discounted due to the limitation of a string matching syntactic approach [3]. A semantic approach was chosen because of the need to understand the complex relationships between system capabilities. The design focus on ontology artifacts in relation to the whole discovery process is the central theme — exploring search performance with practical Financial Services (FS) knowledge topologies.

The remainder of this paper is organized as follows. Section 2 covers related work. Section 3 shows the financial service architecture. In section 4 the implementation of the prototype is described. Section 5 contains the performance analysis. A conclusion is given in section 6 which summarizes the findings.

## 2. Related Work

Many strategies have been proposed to address the issues of adapting to new distributed business requirements, either through the design of

knowledge flows as seen in Nonaka's knowledge spiral [4] or the technology focus on explicit workflow design (BPEL4WS [5]). Both approaches do not directly address an environment of changing business requirement and strategy across business silos. The explicit knowledge created from experience is valuable, but not in a system usable format and as such cannot be readily utilized when searching for existing capability electronically. Alternatively, the workflow approaches provide an explicit definition of how a capability is used at both a macro (the workflow) and micro level (the service). One problem here is that no single person, or pre-formed grouping, can bridge the wide conceptual knowledge gap between the macro and micro perspectives. Chen et al. [6] add, "such initiatives (WSFL, XLANG, BPEL4WS) generally focus on representing service compositions where the flow of process and the bindings between services are known a priori. . .".

Similar research in the Grid area was addressed by Deelman et al. [7] with their workflow generator and Tangmurarunkit et al. [8] with their resource selector. The workflow generator addresses the problem of automatically generating job workflows for the Grid. The system operates in the application domain and chooses application components based on the application metadata attributes.

The ontology-based resource selector exploits ontologies, background knowledge, and rules for solving resource matching in the Grid to overcome the restrictions and constraints of resource descriptions in the Grid. The user can activate the matchmaker by submitting a query asking for resources that satisfy the request specification. The query is then processed by the TRIPLE/XSB deductive database system [9] using matchmaking rules, in combination with background knowledge and ontologies to find the best match for the request. The use of richer explicit knowledge, that is system readable, provides the basis for discovering capability on the Business Grid — "the grid of services" [10]. This diversity of human and system focus is bridged within this research by focusing on the use of multiple, co-existing service ontologies as opposed to more centralized or homogenized resource centric approaches.

### 3. Financial Service Architecture

Business and technical architecture of financial service applications may be segmented by product, process or geographic concerns. Segmented inventories make inter silo re-use difficult. The brittle nature of application integration is then replaced with that of service integration. An alternative approach is to adopt a loosely coupled inventory — containing differing explicit capability knowledge.

In order to investigate the practical use of capability knowledge capabilities from a group of systems that spanned the product, process and geographical dimensions were extracted. The capabilities were explicitly described using OWL (W3C Web Ontology Language [11]) — see section 3 — and exposed as Grid services. Three use-cases were chosen to explore the use of semantic searching:

- *Use-case 1 — Searching for trades executed with a particular counterparty*
- *Use-case 2 — Valuing a portfolio of interest rate derivative products*
- *Use-case 3 — Valuing an option-based product*

These three use-cases were chosen because they provide examples of three distinct patterns of use — aggregation, standard selection and multiple selection.

The Semantic Discovery for Grid Services architecture (SEDI4G) (depicted in Fig. 1) allows placement of the discovery services across

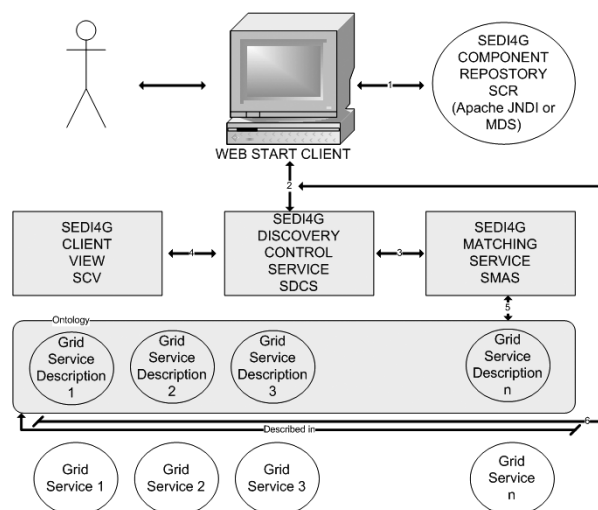


Fig. 1. SEDI4G Architecture.

the network. Several applications from this domain were mined for capabilities (e.g. methods and parameters). These capabilities were then stubbed as Grid and web services with examples such as methods for the trade pricing of Interest Rate Swap and Interest Rate Option products resident in both trading and risk management systems. The use of *de facto* business terminology and object generalization can result in certain terms requiring additional analysis. An example is the use of the term “trade” to describe the various states of the trade — planned, new, live and dead. These specializations were recognized and made explicit when describing the services in ontological form as this gave greater selection accuracy.

The discovery process begins by identifying components required to carry out a semantic search. This choice involves mainly placement concerns (represented by the grey flexible services and data in Fig. 1). Thus, Step 1 involves the selection of which discovery control service (SDCS), knowledge base and matching service best fit the user requirement. This information is sent to SDCS together with the search parameters (2). SDCS then calls the KB-based matching service (3) that in turn loads the KB and rules (5). The matching is carried out and returned to SDCS for use in one of the client components (4). The SDCS service can optionally provide the resource properties, the dynamic state of each service, alongside the service choices (6).

As the matching algorithm is based on a semantic method using an ontology and a reasoning engine, the background to the subject is introduced. The subject of ontology, a term borrowed from philosophy [12], is the study of the categories of things that exist, or may exist, in some domain [13]. An ontology is a catalogue of the types of things that are assumed to exist in a domain of interest from the perspective of a person who uses a language for the purpose of talking about a domain. The types in the ontology represent the predicates, word meanings, or concept and relation types of the language when used to discuss topics in the domain. An uninterpreted logic is ontologically neutral: It imposes no constraints on the subject matter or the way the subject is characterized. Logic alone says nothing about anything, but the combination of logic with an ontology provides a language that can express relationships about the entities in the domain of interest.

The matching algorithm comprises two steps; the initialization of the knowledge base and the actual search request. During the initialization phase the ontology is loaded and the facts and rules are applied using the Rete algorithm [14]. During the search request the ontology is reasoned depending on the specified Jess queries and the matched services are returned.

#### 4. Implementation

A prototype system was implemented using Protégé (with OWL Plugin) to build three service KBs — OWL files (small, medium and large). The discovery components are Grid-enabled web services. The control web service is a Java web service that directs control to the discovery algorithm which then generates either; (a) XML for return to the client, (b) HTML for thin client usage and (c) a Java Desktop JDNC data file. Research into distributed architecture poses many changes; elimination of middleware overhead and suitable tooling being examples. Granular logging allowed for the removal of such overhead and tools with flexible code placement were chosen in support of configuration flexibility. Realistic search use-cases, utilizing services of differing levels of specialization, directed the chosen search tests (see section 5).

To provide an enhanced discovery for the FSs, the prototype is supported by a matching mechanism which allows for a more efficient service discovery by using a FS ontology described in a semantic language and a reasoning engine that uses the ontology [15].

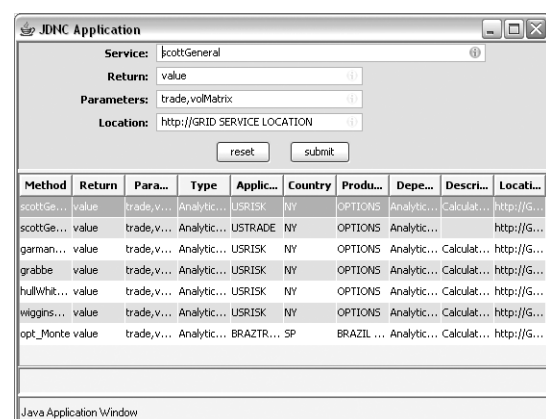


Fig. 2. Service Search and Selection.



Fig. 3. FS Ontology Structure.

OWLJessKB [16] java library, used within the SMAS service, was used to read and query the OWL files. It inserts the RDF triples as facts into the Jess knowledge base [17]. With some predefined rules, Jess can reason about the triples and can draw more inferences. The Jess API (Application Programming Interface) is intended to facilitate interpretation of information of OWL files, and it allows users to query on that information. It leverages the existing RDF API to read in the OWL file as a collection of RDF triples. How the service discovery process works is shown next.

Fig. 3 shows a part of the FS Ontology. The services are listed as classes of the ontology, having methods, which are sub-classes, and parameters which are properties spoken in ontology terms. The user defines for example three search words which are: `executeTrade`, `boolean` and `trade`. For all three search words the Jess query shown in Fig. 4 is invoked. The returned classes are then taken and the Jess query shown in Fig. 5 is invoked for each return class. This returns the services `getTradeByID`, `getTradeByDate`, `getTradeByType` etc. The sub-classing query finds the other three related services, which is basi-

```
(defquery query-for-class-of-a-given-property
  "Find the class to a given property.")

(declare (variables ?class))

(triple
  (predicate 'http://www.w3.org/2000/01/
             rdf-schema#domain')
  (subject ?class)
  (object ?x)
)
)
```

Fig. 4. Jess Query — Property of Class.

```
(defquery query-for-class
  "Find all the sub-classes.")

(declare (variables ?y))

(triple
  (predicate 'http://www.w3.org/2000/01/
             rdf-schema#subClassOf')
  (subject ?x)
  (object ?y)
)
)
```

Fig. 5. Jess Query — Sub-Classing.

cally done with the subsumption relation which defines a hierarchy of concepts.

Returned services are then displayed in the GUI (Fig. 2).

## 5. Performance Measurements

The prototype system used three ontologies (summarized in Table 1) to investigate the impact of KB size on semantic search performance. The ontology growth will be driven by increases in the number of product, process or geographical systems being described. The following setup was chosen. The three use-cases mentioned in section 2 had the following query attributes:

- *Use-case 1: Finding a trade executed against a particular counterpart* [Query= "tradeList", "cpt"]
- *Use-case 2: Finding a capability to value a book of interest rate products* [Query= "rateVector", "curveStructure", "doneTrade"]
- *Use-case 3: Finding a capability to value an option trade* [Query= "liveSwaptionTrade", "volMatrix"]

Ontology Size	Small	Medium	Large
No. of classes	57	114	228
No. of prop.	23	46	92

Table 1. Ontology Comparison.

In use-case 1, the user has a list of trades executed with a particular counterparty. The ontological model identifies four services — of type credit risk analysis and portfolio persistence. Specific process attributes could further refine the result set to risk or persistence services.

Two sets of measurements were taken — ontology size (Fig. 6) and placement (Fig. 7).

The timings in milliseconds cover server time, both in instantiating the OWLJessKB objects and executing the semantic search. Client side and initialization figures highlight the need for server side caching as initial runs reveal an overhead in the region of 10 seconds. The timing labels are new object initialization (N), finding semantic matches (M) and the total (T).




Fig. 6. Performance Test — Local.




Fig. 7. Performance Test — Local vs. Distributed.

Ontology Size	Perf. of N in MS	Perf. of M in ms	Perf. of T in ms
Small	196	1158	1354
Medium	188	2255	2443
Large	181	5171	5353

Table 2. Average Result of Local Performance.

They are taken from services that have been initialized. As seen from the measurements, the longest timings occur during the matchmaking process, due mainly to the Rete algorithm.

Table 2 summarizes the average results of the first set of measurements. The measurements taken show an average variation of 8ms, 67ms and 63ms of the time taken for N, M and T respectively. This is as expected as the time for N is relatively small in comparison to M and T.

Fig. 7 shows the performance measurements of three use-cases using local and distributed KBs. Comparing the results quantitatively reveals that the communication overhead has a bigger impact on the small ontology than on the large ontology. Measured for this overhead were 24%, 28% and 24% for the small ontology and 5%, 4% and 3% for the large ontology. This is due to the fact that the time taken for loading the small ontology and the application of the rules were much smaller than for the larger ontology.

Table 3 summarizes the average performance values for the local versus distributed setup. This testing was carried out as earlier results highlighted the need for segmented knowledge, distributed over the network. The implications of this communication overhead are that on a good (broadband or better) network the relative overhead is small (4%) when searching several hundred classes. The overhead increases to 25% for the small distributed KB.

Ontology Size	Local Performance	Distributed Performance
Small	883ms	1182ms
Large	5362ms	5576ms

Table 3. Average Local vs. Distributed Performance.

## 6. Conclusion

This resulting system and performance results show that dynamic re-use of service capabilities is easily supported by the discovery service. The performance of the semantic search (when fully optimized) is acceptable for use within this domain for ontologies up to 200-300 classes (the large ontology test case). The performance grows linearly with the number of classes. Likely, ontology sizes within an Investment Bank extend to many thousands of classes directing further work on specialized KBs and other optimization methods. Distributed KB support is a key characteristic of the SEDI4G architecture, allowing for many KBs to exist within its search scope (a prerequisite when thousands of classes exist across the organization). This distribution, on a performant network, has limited impact when dealing with KBs containing hundreds of classes (exposed in Table 3); and this should be recognized when architecting such topologies.

## References

- [1] UDDI Technical White Paper. <http://uddi.org/pubs/uddi-tech-wp.pdf>.
- [2] S. FITZGERALD, I. FOSTER, C. KESSELMAN, G. VON LASZEWSKI, W. SMITH AND S. TUECKE, A Directory Service for Configuring High-Performance Distributed Computations. *Proceedings of the 6<sup>th</sup> IEEE International Symposium on High-Performance Distributed Computing* (HPDC-6), 1997.
- [3] L. LI AND I. HORROCKS, A Software Framework for Matchmaking Based on Semantic Web Technology, *International Journal of Electronic Commerce*, 8 (4). pp. 39–60.
- [4] I. NONAKA, The Knowledge-Creating Company, *Harvard Business Review*, 69 (6). 96, 1995.
- [5] BPEL4WS - Business Process Execution Language for Web Services Version. <http://www-128.ibm.com/developerworks/library/ws-bpel/>.
- [6] L. M. CHEN, N. R. SHADBOLT, C. GOBLE, F. TAO, S. J. COX, C. PULESTON AND P. R. SMART, Towards a Knowledge-based Approach to Semantic Service Composition, *Proceedings of 2nd International Semantic Web Conference* (ISWC2003), 2003.
- [7] E. DEELMAN, J. BLYTHE, Y. GIL, C. KESSELMAN, G. MEHTA, K. VAHI, A. LAZZARINI, A. ARBREE, R. CAVANAUGH, S. KORANDA, Mapping Abstract Complex Workflows onto Grid Environments, *Journal of Grid Computing*, Vol. 1, No. 1, pp. 9–23, 2003.
- [8] H. TANGMUNARUNKIT, S. DECKER, C. KESSELMAN, Ontology-based Resource Matching in the Grid — The Grid Meets the Semantic Web, *Proceedings of the International Semantic Web Conference*, Budapest, Hungary. May 2003.
- [9] The XSB Research Group. <http://xsb.sourceforge.net>.
- [10] F. LEYMAN AND K. GUNTZEL, The Business Grid: Providing Transactional Business Processes via Grid Services, *Proceedings of Service-Oriented Computing* (ICSOC2003), 2003.
- [11] W3C Web Ontology Language. <http://www.w3.org/TR/owl-ref/>.
- [12] T. R. GRUBER, A Translation Approach to Portable Ontology Specifications, *Knowledge Acquisition*, Vol. 5, No. 2, 1993.
- [13] J. F. SOWA, Ontology, Metadata, and Semiotics, Conceptual Structures: *Logical, Linguistic, and Computational Issues*, Lecture Notes in AI #1867, Springer-Verlag, Berlin, pp. 55–81, 2000.
- [14] C. L. FORGY, Rete: A Fast Algorithm for the Many Pattern/Many Object Pattern Match Problem. *Artificial Intelligence*, 19(1982), pp. 17–37.
- [15] S. A. LUDWIG, Flexible Semantic Matchmaking Engine, *Proceedings of 2nd IASTED International Conference on Information and Knowledge Sharing* (IKS), AZ, USA, 2003.
- [16] OWLJessKB. <http://edge.cs.drexel.edu/assemblies/software/owljesskb/>.
- [17] JESS, Java Expert Systems Shell. <http://herzberg.ca.sandia.gov/jess/docs/61/index.html>.

Received: June, 2005.

Accepted: October, 2005.

Contact address:

David Bell  
Department of Information Systems and Computing  
Brunel University, UK  
David.Bell@brunel.ac.uk

Simone A. Ludwig  
School of Computer Science  
Cardiff University, UK  
Simone.Ludwig@cs.cardiff.ac.uk

---

DAVID BELL is a researcher at Brunel University, undertaking a PhD. He holds a MBA from Warwick University, UK and a BSc(Hons) in Computer Science and Electronics. Prior to that, David has worked extensively in industry, with over 10 years experience in investment banking technology. His research interests include grid computing, distributed systems, semantic technology and service orientation. David is a student member of the IEEE, ACM and UKAIS.

---

SIMONE A. LUDWIG is a Research Associate at Cardiff University, UK. She received a Ph.D. degree in Computer Science and an M.Sc. degree in Distributed Computing Systems Engineering from Brunel University, UK in 2004 and 2000 respectively. Her research interests include grid computing, distributed computing, semantic Grid/Web, knowledge engineering and software development. Prior to that, she worked several years as a software engineer for Daimler-Chrysler, Stuttgart, Germany. She is a member of the IEEE.

---