UDDI for a manufactured product brokering service

G. Cope
Peace Software Ltd, Auckland, New Zealand

R.W. Amor
Department of Computer Science, University of Auckland, Auckland, New Zealand

ABSTRACT: The Universal Description, Discovery and Integration (UDDI) initiative provides for distributed web-based information registries of web services. UDDI registries can be used to promote and discover distributed web services. This provides an alternate model to that offered by catalog vendors to date where each manufacturer has to provide separate information to each catalog within which they wish to list their products. This project describes an implementation of UDDI registries to test their ability to handle product manufacturers and suppliers in a brokering framework. It provides a uniform interface for manufacturers and suppliers to define their business offerings and their ability to provide further information to interested parties. The system also provides a product search interface for users to identify products matching a particular classification and discusses how this could be extended to include product parameters.

1 INTRODUCTION

Many research projects, and current catalog providers, have attempted to introduce IT-based solutions to the problem of identifying manufactured products. The majority of these projects appear to have met with little commercial success, for which there are a variety of identified barriers (Jain & Augenbroe 2000). In this project a new technology, which takes advantage of the open nature of the Internet, is tested to ascertain its relative merits for a manufactured product brokering service.

1.1 *Current Manufactured Product Services*

Manufacturers and suppliers publish a huge range of information on their individual products and services. Often these product sheets contain information vital to the decisions on planned functionality of a building and choice of product to match specified functionality. In most western countries there are commercial organizations which collate a large range of product information in the form of a catalog to be on-sold to design firms. These paper catalogs tend to be very large, costly, and being based upon a large collection of independent manufacturers and suppliers soon become out of date. Finding information in these catalogs is usually driven by a particular classification system for manufactured products which just identifies all manufacturers and suppliers of a type of product. Jain & Augenbroe (2000)

discuss how product information is viewed and used in projects.

To address the currency issues of this information, and to enable a more sophisticated search, electronic product catalogs have been trialled. However, most are almost direct copies of their paper counterparts and hence provide little more functionality in their product search capabilities, though potentially improving on their information currency. A few catalog producers have looked at the use of product parameters within their offerings to provide sophisticated and directed product identification, but currently do not offer this ability in their commercial services (Barbour 2000).

A few research projects have developed frameworks for product identification through parameters. The ARROW project (Newnham & Amor 1998) and the MPS service within CONNET (Turk & Amor 2000) both created frameworks which could accept manufacturers' and suppliers' product information in a parameterized form and hence allow users to identify products well suited to their project's needs. The difficulty of collation of parameterized information from manufacturers and suppliers was one of the main barriers to the further uptake of these services.

1.2 *A New Brokering Approach*

The approach taken in this project is to utilize new Internet techniques to support the following functionalities:

- Give manufacturers and suppliers the right, and responsibility, to maintain their own information. As the owners of the information used in the industry they should be able to decide what is made available, when new information is made available, and when modified information is announced.
- Create a system based upon recognized standards so that manufacturers and suppliers do not have to supply almost the same information to many catalog companies in varying formats and styles. Manufacturers and suppliers should be able to provide their information in one format and have it picked up by any of the catalog or brokering companies they choose.
- Work within a framework where information only has to be entered once and can then be disseminated to everyone who subscribes to it, whether they are end users of the information or intermediary catalog producers or brokers.
- Allow manufacturers and suppliers to tie their information offerings through to their local database or information management system. This would allow them to retain control of, and responsibility for, the information that is disseminated on their products.
- Provide users in the industry with an effective mechanism for identifying products and retrieving information available about a product. This mechanism should be accessible to software tools to allow automated interrogation and downloading of information by standard software applications used within the industry (e.g. within CAD systems).

Initial publicity about UDDI suggests that this initiative would allow these modes of operation to be supported for product information.

2 UDDI

On September 6, 2000, a group of technology and business leaders announced that they had come together to develop the Universal Description, Discovery and Integration (UDDI) specification, an initiative to create a global, platform-independent, open framework to help achieve the following goals (UDDI 2001):
- Enable businesses to discover the right trading partners out of all those available online.
- Define how to conduct business once the right partners are located.
- Create an industry-wide approach for businesses to quickly and easily integrate with their customers and partners on the Internet with information about their products and services.
- Describe how businesses prefer to be integrated into each other's systems and business processes.

The UDDI project is an industry initiative lead by Accenture, Ariba, Commerce One, Compaq, Fujitsu, Hewlett-Packard, i2 Technologies, IBM, Intel, Microsoft, Oracle, SAP, Sun Microsystems, and Verisign. More than 300 other companies participate in the UDDI advisors group.

The UDDI project is an attempted comprehensive solution for businesses to publish their information to any customer or business partner around the world. This is achieved by the creation of a services registry architecture that presents a standard way for businesses to build a registry, query other businesses, and enable those registered businesses to interoperate and share information globally in a distributed manner.

The UDDI framework allows businesses to describe the business services their web sites offer, and how they can be accessed globally over the web. It defines standards for how businesses will share information, what information they need to make public, what information they choose to keep private, and how to describe their businesses and services. It reaches far beyond today's Internet business listings and search directories that provide specific but limited value to an organization.

The UDDI specifications take advantage of World Wide Web Consortium (W3C) and Internet Engineering Task Force (IETF) standards such as Extensible Markup Language (XML), HTTP, and Domain Name System (DNS) protocols. Additionally cross platform features are addressed by use of Simple Object Access Protocol (SOAP) messaging specifications.

UDDI is a general-purpose registry. It can be used to register any type of service, not just web services. UDDI is really two things:
1. The UDDI cloud of operator nodes, an Internet-wide repository (made up of white, yellow, and green pages) for web services metadata. Figure 1 provides a description of the types of information that are offered through the UDDI repository.
2. An API and data model standard for a web services metadata repository. The latter hosts the data while the former provides a means to access it (Graham 2001).
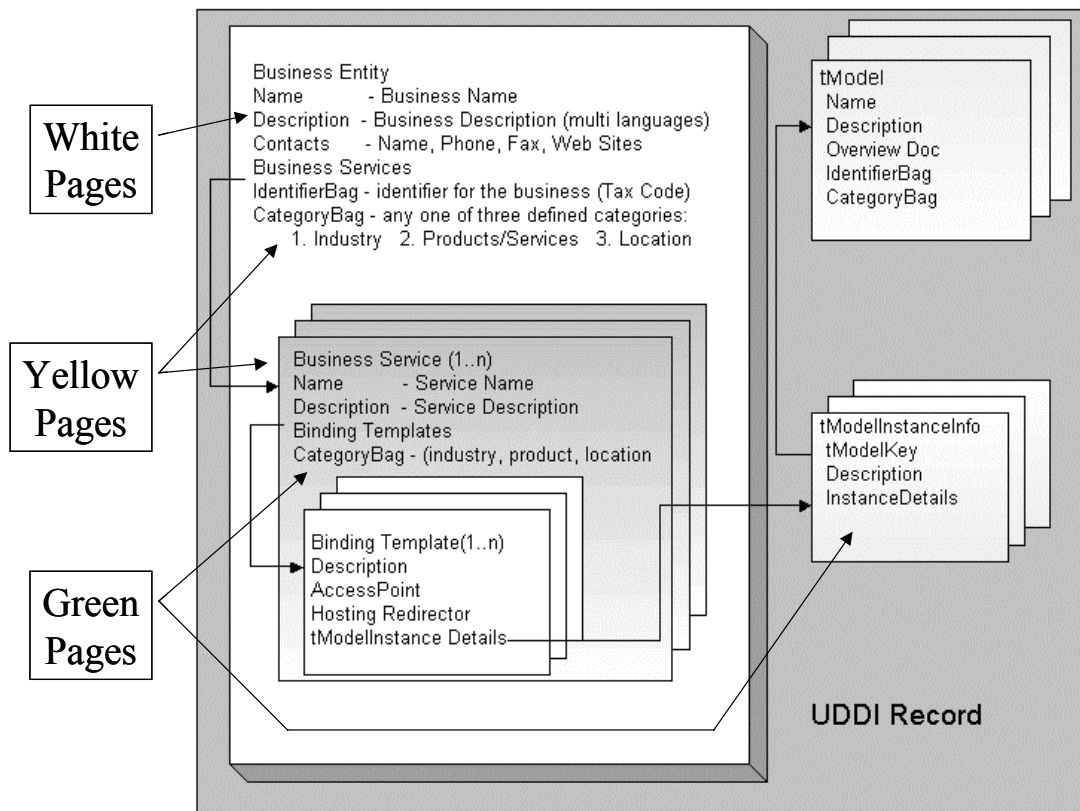
Figure 1. UDDI core data structures (Siddiqui 2001)

2.1 *UDDI Data Structures*

UDDI defines four basic data structures within the data model: Business Entity (modeling business information), Business Service (high level service description), tModel (modeling a technology type or service type), and Binding Template (mapping between Business Service and tModels). Within these basic data structures there are also important attributes for identification of Business Entities and tModels, and classification of Business Entities, Business Services, and tModels. Figure 2 provides a description of the relationships between the data structures.

2.1.1 *Business Entity*

The business entity data structure is the principal entity within UDDI.

Within the business entity are the core details of the business: its name, descriptions of the business, details of personal contacts within the business, identifiers (identifier bag) used to refer to the business (e.g. DUNS, IRD Number, etc), categorical classifications (category bag) for the business (e.g. NAICS or ISO 3166 for geographical location), and a business services structure which defines the services of the business. The business entity's name and contact details can be typically seen at as the white pages information of the UDDI registry.

2.1.2 *Business Service*

Each business service data structure is associated with a single business entity and describes a single service offered by the business.

A service consists of a name, a description, a classification of the service type (category bag), and a binding templates structure which holds the technical service description information related to a given business service family. The name and description may be repeated in other languages for international usage.

The business service, along with category bag of both the service and its associated business entity, can be seen as the yellow pages information of the UDDI registry.

2.1.3 *Binding Template*

Each business service will contain a number of binding templates. These binding templates represent the technical information of "where" a business service may be contacted. This will typically be the entry point address(es) of a web service. Contained within each binding template will be references to instances of tModels, the information required to know how to do business with this service.
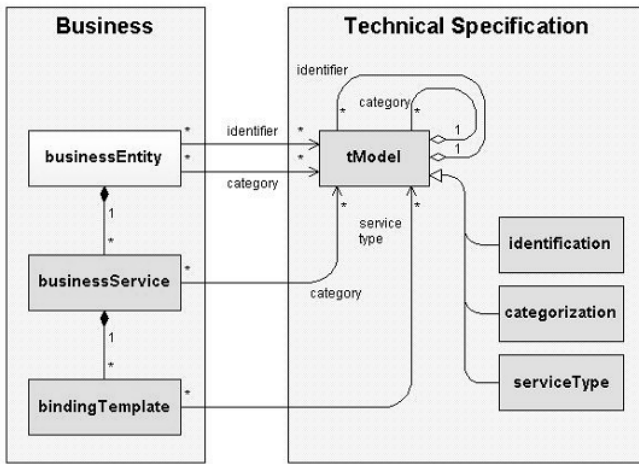
Figure 2. Relationships between UDDI data structures (von Riegen 2001)

Binding template information will typically be retrieved by an organization interested in using the related business service. This particular information can then be used for dynamic invocation at runtime. Used in this manner a publishing organization can use the UDDI registry as a way to suddenly change its service offerings and entry point without having to notify all its business partners of this change.

Together the binding template and the tModel form what can be seen as "green pages" (or technical information) about a business's services.

2.1.4 *tModel*

A tModel, or technical model, represents two things within UDDI: a technical specification for a given service type, or a model for a particular identifier or taxonomy. Examples of tModels would include WSDL files, XML DTDs or schemas, and classification schemes.

To describe the technical specification of the connection to a service type, where multiple businesses may offer the same interface to their service, a reference to a unique tModel would be used. If a system knows how to interact with a specification referenced by a particular tModel it can then identify which services it can work with dynamically.

A tModel would specify information such as its name, the name of the publishing organization, a list of categories that identify the service type, and pointers to technical specifications (e.g. interface definitions, message formats, and message and security protocols).

An example of a tModel for a taxonomy would be one for the use of the CSI classification. Any time a business service is seen with the tModel for CSI then an application which understands CSI would know that it can use the associated classification code.

2.1.5 *Identifier Bag*

The identifier bag element allows business entity or tModel structures to include information about common forms of identification (e.g. DUNS, IRD Number etc). This can be used to signify the identity of the business entity or that of the publisher. When used it greatly enhances the search behavior exposed by find_xx messages used with UDDI (2001).

2.1.6 *Category Bag*

The category bag element allows business entity, business service and tModel structures to be categorized according to any available classification schemes. The UDDI public registries provide validated support for a number of taxonomies that cover industry codes (NAICS), product and service classifications (UNSPSC), and geography (ISO 3166). Domain specific classifications can be specified and included in the category bag (e.g. CSI, UNICLASS, EPIC, CBI, etc). Inclusion of classifications is optional but again greatly enhances search behavior exposed by find_xx messages (UDDI 2001).

3 THE MANUFACTURED PRODUCT BROKERING SERVICE

The manufactured product brokering service developed in this project has a structure as shown in Figure 3. A private UDDI registry was created to contain information on manufacturers and suppliers of interest to the construction industry. This private registry is able to push business information through to other open UDDI registries, however it does not allow non-validated information to be propagated back. The broker and client services are described in the following sub-sections.
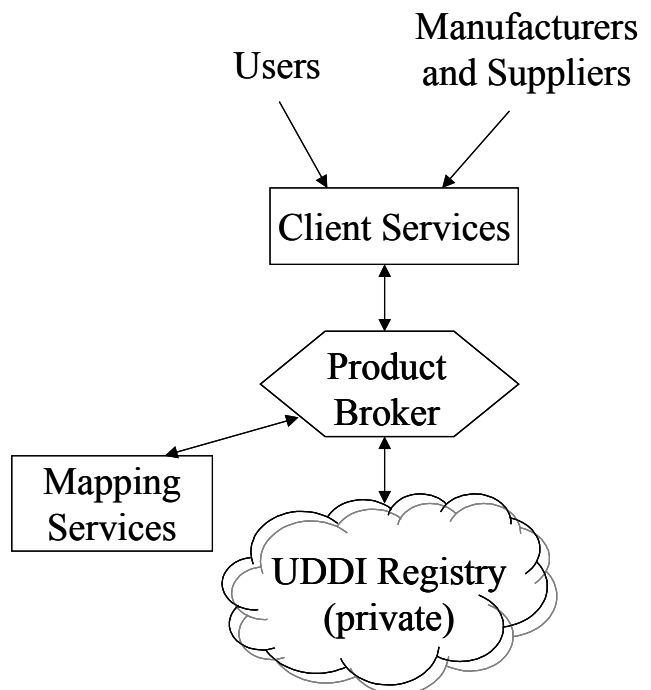

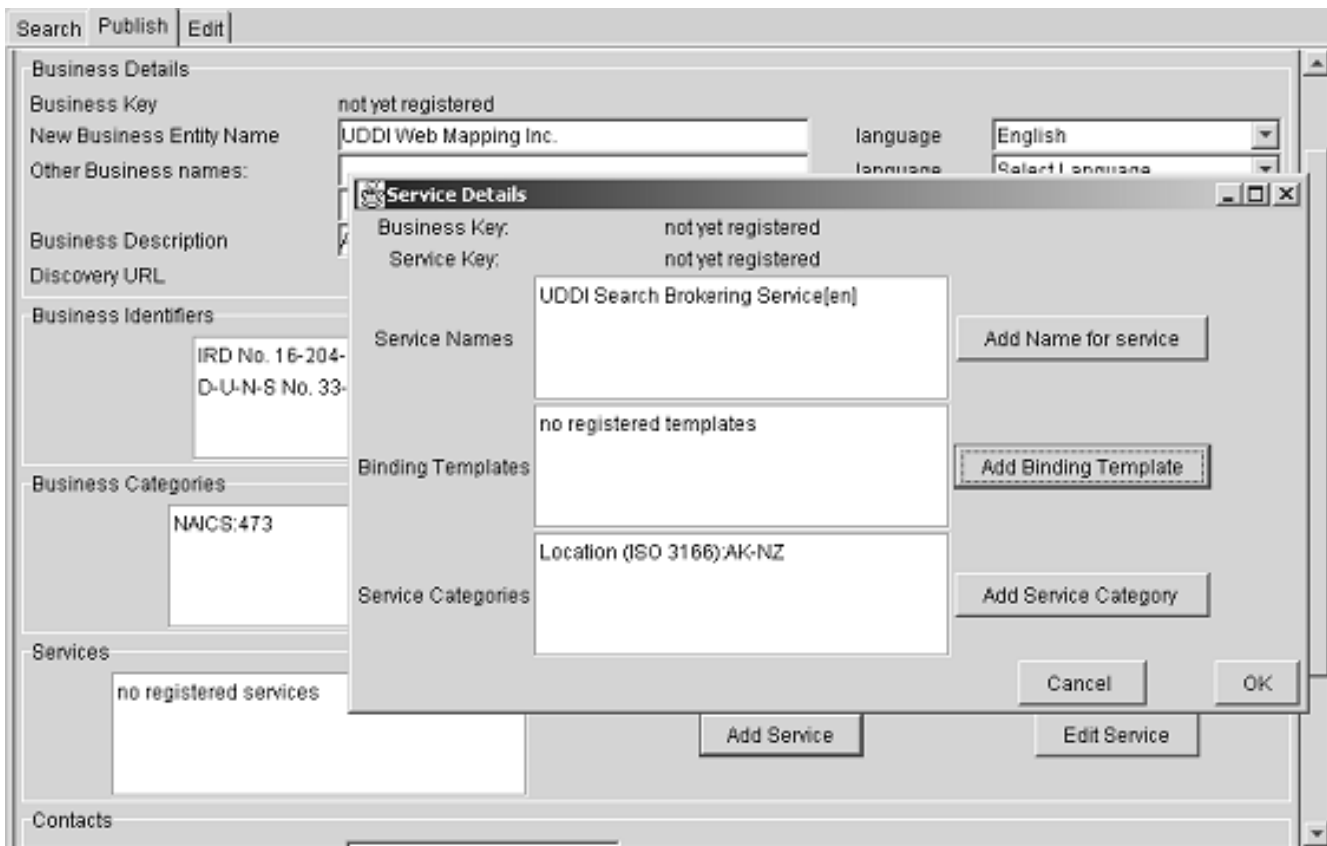Figure 3. Structure of the product broker

Figure 4. Publishing business information

3.1 *Broker services*

The brokering service provides bespoke functionality for construction industry requirements along with the standard UDDI functionalities as follows:

− Providing clients with the ability to register, and validate, themselves to the UDDI registry.
− Enabling clients to publish, and edit, the UDDI standard and construction industry specific business information in the UDDI registry. The broker ensures that the information in the UDDI registry is valid and complete to the requirements specified for manufactured product brokering.
− Providing a search capability on top of the UDDI registry against UDDI standard and construction industry specific information for businesses and services which match particular requirements.
− Providing classification specific translation functions for a range of industry specific classification systems. This enables manufacturers and suppliers to use their normal classification system when describing their services. The broker provides translation facilities in order to match requests against the provided classified information.

The management of classifications, other than those built into the UDDI system or supplied via a tModel specification, proved somewhat problematic. The UDDI system allows for classifications other than UNSPSC, NAICS, etc but all non-supported ones are tagged as 'other' within the category bag. While this does not impede the storage of other classification codes (other than forcing the broker to validate the codes) it does hinder the selection of businesses and services by these classification codes. Basically, all 'other' classification codes must be searched sequentially to identify whether a specified code exists in a category bag or not. For a system with a very large number of products and services this will prove to be very slow. For optimal performance UDDI-based brokering will require the use of the built in and generic classification schemes, or the provision of a wide range of construction industry specific taxonomies through the tModel mechanism.

3.2 *Client services*

The client services developed for the manufactured product broker provide user interfaces and APIs for the range of functions typical clients would wish to undertake. These include:

− Providing manufacturers and suppliers with the interfaces required to publish, and edit, business and service information. Figure 4 shows a user interface to enter service details for a business.
− Providing users with the interfaces required to search for businesses, or services, or products. This search can be based on any of the information within the UDDI registry. This

includes business names, service and product names, classification codes for product types, locality of a business service, etc.

– Providing users with the ability to navigate through the UDDI data structures to find information related to their search queries.

In providing the client functionality it was clear that a major aspect of the information that users require about manufactured products (i.e. product parameters) could not be easily incorporated into the UDDI registry system. UDDI is targeted at establishing business relationships and assumes that the business itself provides all further information about its offerings. So while identifying a product supplied by a business is within scope, enabling the user to find a product based on functional parameters is outside the UDDI registry's scope.

UDDI does, however, provide a series of mechanisms which would make this level of functionality achievable. The UDDI binding template provides information on how to interact with the specified business and service. This can take the form of WSDL (Web Services Description Language) and WSCL (Web Services Conversation Language) which describe the data structures, procedure calls, and process that needs to be followed to achieve a functionality from a business' service.

This pushes the onus for providing product information back to the manufacturer and supplier. With the huge numbers of manufacturers and suppliers in this industry, and their current level of IT capability, this will ensure that very little of this information will come through. It would appear that catalog providers, with their current business models, could be a conduit for the provision of parameterized information as required for this level of product brokering.

4 CONCLUSIONS AND FUTURE WORK

This project created a UDDI-based service for the brokering of manufactured products. UDDI was selected to trial a new model of brokering. This model provides manufacturers and suppliers with the ability, and responsibility, to manage their business and service related information in a standard format through a single registry, but for this information to be available to any number of services which may wish to build on it (e.g. catalog-based enterprises). This portion of the project was successful.

However, it has become clear that UDDI is not set up to allow parameterized information to be included about products and services. This is a serious restriction on the ability to create a fully functional manufactured product brokering service, as it requires all manufacturers and suppliers to publish the parameters of their products individually.

Given the number of manufacturers and suppliers and their general IT capabilities this is not currently feasible.

Future work on this system aims to address the more problematic issues of managing parameterized information between the broker and manufacturers and suppliers. We aim to establish a UDDI like registry for parameterized information that could be connected through the current broker. We also aim to shift the major construction related classifications into the UDDI recognized set. This will enable more efficient retrieval of construction related information from UDDI based registries.

REFERENCES

Barbour. 2000. *Barbour index building product expert*, web site last accessed 5/5/2002, http://www.barbourexpert.com/.

Graham, S. 2001. *The role of private UDDI nodes in Web services*, web site last accessed 5/5/2002, http://www-106.ibm.com/developerworks/webservices/library/ws-rpu1.html?dwzone=webservices.

Jain, S. & Augenbroe, G. 2000. The Role of Electronic Product Data Catalogs in Design Management, *CIB W96 Conference on Architectural Management*, May, 271-288.

Newnham, L. & Amor, R. 1998. Translation of Manufacturer's Product Data for the ARROW Product Search System, *the second European Conference on Product and Process Modelling*, London, UK, 19-21 October, 405-412.

Siddiqui, B. 2001. *Using UDDI as a Search Engine: smart web crawlers for all*, web site last accessed 5/5/2002, http://www.webservicesarchitect.com/content/articles/siddiqui01.asp

Turk, Z. & Amor, R. 2000. Architectural foundations of a construction information network, *International Journal of Construction Information Technology*, 7(2): 85-97.

UDDI. 2001. *UDDI Executive White Paper*, November 14, web site last accessed 5/5/2002, http://uddi.org/.

UDDI. 2002. *Universal Description, Discovery and Integration of Business for the Web*, Version 2.0 Specifications, web site last accessed 5/5/2002, http://uddi.org/.

von Riegen, C. 2001. Connecting Trading Partners and Marketplaces with UDDI, *XML Europe 2001*, Berlin, Germany, 21-25 May, web site last accessed 5/5/2002, http://www.gca.org/papers/xmleurope2001/papers/html/s26-2.html