PERSONALISED RESOURCE DISCOVERY SEARCHING OVER MULTIPLE REPOSITORY TYPES

Using user and information provider profiling

Boris Rousseau, Parisch Browne, Paul Malone, Mícheál ÓFoghlú
Telecommunications Software Systems Group (TSSG), Waterford Institute of Technology, Waterford, IRELAND

Paul Foster, Venura Mendis
BTexact Technologies, Orion Building, Adastral Park, Martlesham Heath, Suffolk, IP5 3RE, United Kingdom

Keywords: E-Learning, personalisation, profiling, user profile, information provider profile, resource discovery, multiple repository types.

Abstract: The success of the Information Society, with the overabundance of online multimedia information, has become an obstacle for users to discover pertinent resources. For those users, the key is the refinement of resource discovery as the choice and complexity of available online content continues to grow. The work presented in this paper will address this issue by representing complex extensible user and information provider profiles and content metadata using XML and the provision of a middle canonical language to aid in learner-to-content matching, independent of the underlying metadata format. This approach can provide a federated search solution leading to personalise resource discovery based on user requirements and preferences, seamlessly searching over multiple repository types. The novelty of the work includes the complex extensible user profiles, information provider profiles, the canonical language and the federated search strategy. Although, the work presented is focused on E-Learning, the general ideas could be applied to any resource discovery or information retrieval system.

1 INTRODUCTION

Most current E-Learning resource discovery systems (GESTALT, 1998) employ relatively simple algorithms when searching for content. The more complex of these resource discovery services (RDS) may use a thesaurus to find similar words and build a set of related pages. This allows the user to be provided with a 'Find Similar Pages' option. Although early research in the area of Artificial Intelligence (Spark, 1978) provided some solutions, a successful Information Retrieval system should take into account the identity of the end user performing the search. Such identity incorporates their needs, what type of content they prefer or indeed whether they have available resources to access the content found. Future successful RDSs must consider these factors when executing search queries. This can be best achieved through the use of metadata. Metadata (Day, 2000) is traditionally defined as information on the resources, not only to indicate the nature of the content, but also to provide information about the user as well as the provider of the content.

In this work, models for describing this metadata are explored and new ones developed where necessary. This paper will also demonstrate how this metadata can be used to provide personalised resource discovery, thus facilitating the user's requirement for pertinent information, and the provider's requirement for maintaining multiple, diverse, repository types, while at the same time facilitating the business model that can support this.

The Information Society Technologies (IST) project GUARDIANS (GUARDIANS, 2000) has undertaken this approach and this paper describes a large part of the work. This paper will present and address the issues faced during the project regarding

Rousseau B., Browne P., Malone P., ÓFoghlú M., Foster P. and Mendis V. (2004).
PERSONALISED RESOURCE DISCOVERY SEARCHING OVER MULTIPLE REPOSITORY TYPES - Using user and information provider profiling.
In *Proceedings of the Sixth International Conference on Enterprise Information Systems*, pages 35-43
DOI: 10.5220/0002649800350043

user and information provider profiling as well as the search strategies to enable customisable resource discovery. The first section will cover the user and information provider profiling, in addition to content metadata extensions for preferences weighting facilitating the ranking of results found. The main section will present the work carried out on the canonical language and the federated search strategy handling multiple repository types to provide personalised resource discovery.

2 PROFILING

2.1 User Profile

With the growth of e-commerce and the Internet, the gathering of information on preferences, activities and characteristics of clients and customers has become quite important, as this information can facilitate personalised content delivery. The idea of user profiling can be summarised as "one that renders the user with an experience that is tailored to his/her current situation" (Suryanarayana, Hjelm 2002). The issue of user profiling is quite complex as such profile should be as generic as possible with a compromise between storing too little and too much information. A starting point is to use XML (eXtensible Markup Language, 2000), in order to be exportable to other system components and to incorporate information about the user's preferred language and the technological platform for communication.

As described in (Rousseau et al., 2003), investigation of the existing profiling standards revealed that while each of the standards investigated had its own merits it was felt that none fully addressed the profiling needs of the project. Specifications researched include the IMS Learner Information Package (IMS LIP, 2001) specification, the IEEE Public and Private Information (IEEE PAPI, 2002) Learner standard, and the TV Anytime (TV Anytime, 1999) and vCard (vCard, 1996) Specifications. Based on this research, a complex, extensible, generic user profile model is developed mainly based on the IMS LIP specification, the Generic User Profile (GUP). The profile contains information on the user stored in 'sections' as shown in Table 1.

Table 1: Sections within the GUARDIANS Generic User Profile

GUP Section	Description
Accessibility	User platform, applications, disabilities, languages etc.
Affiliation	Organisations with which the learner has associations or membership.
Usage History	What the learner has previously searched and/or accessed.
Relationship	Relationships between GUP elements.
Interest	Interests of the learner, classified by domain.
Contact Details	The learner's contact details.
Qualifications	Formal qualifications that the learner has previously acquired.
Goals	Learning objectives of the learner.

Although the IMS LIP specification is quite complete, it was felt that extensions where needed to facilitate the search process.

A first extension is the usage history building on the TV Anytime forum specification. Such extension allows to record previous searches information and to reuse it in future ones. A usage history instance fragment follows:

```
<history>
  <InstanceId>history_01</InstanceId>
  <actionhistory>
    <datetime>2002-07-23T20:00</datetime>
    <action>
      <defaulttype>Select</defaulttype>
      <item itemused="INSP">
        <value>NASA</value>
      </item>
      <url>http://www.nasa.com</url>
     </action>
   </actionhistory>
</history>
```

The 'action' subsection depicts whether the user has selected some content (itemused="Content") or an INSP (itemused="INSP"). In the case above, if the user decides to do his next search using "astronomy" as a keyword, "NASA" related content is more likely to come up on his search results.

A second extension allows the user to further refine his/her preferences indicating greater or lesser importance with weightings. These weightings help at the query-building phase to refine further the search and at the rating of results found stages. A weighted fragment follows:

```
<language weighting="10">it</language>
<language weighting="2">en</language>
```

With this in place, the user gets resources preferably in Italian. English content is also suitable but with a lesser importance.

2.2 Information Service Provider Profile

In the standards development community, much of the profiling work seems excessively devoted to user profiles and how they might be used for resource discovery personalisation and communication with a single Service Provider. However, in GUARDIANS this process is seen as only half of the story, the user should choose which service provider he/she wants to enrol with or use. This section details the requirements for storing information on service providers. It will determine what information must be stored in order to adequately describe a service provider and its content and how it may be represented.

The Information Service Provider (INSP) Profile data-model can be used to store information on a service provider on a per-repository basis. It stores both domain specific metadata and generic metadata and defines an aggregated view of the information stored in its repository. As the INSP Profile is essentially an aggregation or summary of the content and services hosted or offered by the INSP, one of its functions must be to act as a collection level description such as those defined by Research Support Library Programme (RSLP, 1999) and UKOLN (UKOLN). In addition it must also function as a searchable resource and be able to operate across multiple domains.

IMS Learning Resource Metadata Information Model Version 1.2 (IMS, 2001) is a specification derived from Version 3.5 Learning Object Metadata Scheme working document of the IEEE LTSC LOM Working Group (IEEE LTSC LOM, 2002). The differences lie in the redefinition of some of the elements, some of which have a new type and/or a new description and/or a new multiplicity. Although both specifications are suitable for the INSP profile, IMS seems to have scope according to up-to-date work coming from IMS and IEEE.

The key information held in the INSP profile is shown in Table 2.

Table 2: Sections within the GUARDIANS INSP profile

INSP Section	Description
Service Name	The common name of the service.
Description	Description of the service.
Contact Details	Information service contact details.
Metametadata	Detail about the metadata contained in the profile.
Legal Constraints	Legal Issues surrounding the use of the content held by the INSP.
Payment information	Details of payment methods accepted by the INSP.
Security Information	Level of security required to access the various elements in this profile.
Relationships	Relationships between instance metadata stored locally or in other INSP profiles.
Classifications	Classification system.
Taxonpaths	A taxonomic path in a specific classification system. Each succeeding level is a refinement in the definition of the higher level.
Keywords	Selected keywords from multiple keyword lists. Keywords can be weighted.

One important element in the INSP profile is the classification. Similarly to a library it represents the categories covered by the service provider. An instance fragment follows:

```
<classification>
  <purpose>
    <choice>Educational content in
    Science</choice>
  </purpose>
  <taxonpath>
    <source>www.cosmos.com</source>
     <taxon>
       <identifier>Stars</identifier>
       <entry>004.41</entry>
     </taxon>
     …
   </taxonpath>
</classification>
```

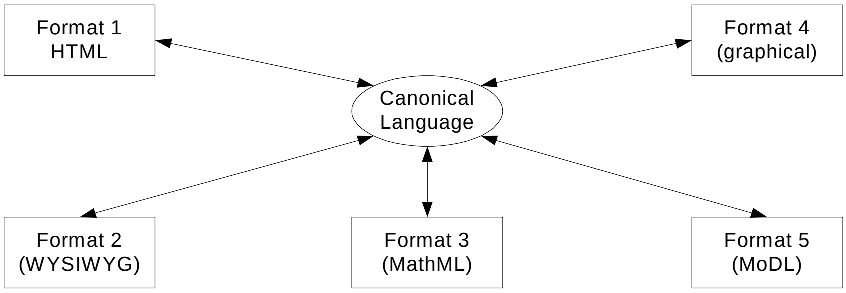

Figure 1: Canonical data model solution

From this extract, the ‘identifier’ section indicates that the service provider offers some resources on the subject of Software Engineering.

There is a requirement in the INSP data model to allow INSPs to offer hosting to content providers in multiple domains or application areas. In addition, some information service providers may wish to extend their profile with other information outside of the standard structure defined above.

The obvious solution is to use a two-category structure: one generic category (for contact details, legal constraints, payment information, etc) and an unbound number of domain specific categories (for keywords and classifications) to handle multiple repository types.

2.3 Storage and Retrieval of profile information

The overall system is based upon profiles. The success of the query depends largely on how those profiles are mapped and on the profile retrieval performance. For storage and retrieval of user and service provider profiles information, different solutions were investigated:

- Local client file system storage.
- Central RDBMS Database storage.
- Central LDAP Directory Service.

Although, storing profile information on a local drive is convenient, the issue of storing space would arise, as the number of profiles gets larger. Such solution would require investigating and implementing compression mechanisms, which would ultimately slow down the profile information retrieval process.

The second solution is to store the data in a relational database system such as Oracle (Chang et al., 2000). The general principles are discussed in (Kanne and Moerkotte, 2000) and (Florescu and Kossmann, 1999). However, at the time of the implementation, no mature, reliable XML database management solution could solve this problem. Also, such a deployment is heavyweight and consequently would affect the systems' performance.

The final and chosen option is to store the data in a Directory Service such as the Lightweight Directory Access Protocol, a simple variant of the X.500 ISO standard (Howes et al., 1995). Some basic classes for persons in an organisation are directly derived from the X500 Directory Services standards. A directory is a specialised database optimised for storage, manipulation and retrieval of information in a globally scalable system, supporting sophisticated filtering capabilities and high volume search request.

However, unlike database management systems, directory services do not support complicated transactions or rollbacks. A profile in a directory service refers to a person or an organisation having one or more relationships with the organisation responsible for the maintenance of the directory. For instance, an Internet Service Provider that stores personal details and access rights. For an LDAP Directory Service to implement a profile, the information should be stored as classes ranging from abstract to more specific ones. For instance, abstract classes could be general information such as contact details and description.

3 CANONICAL LANGUAGE

An underlying concept of the GUARDIANS is that it has no knowledge of the metadata formats on which the repositories are based. Hence, some mechanism is needed to transform the metadata from one format to another. For example, the water molecule (H2O) can be written in HTML as something like:

```
<FONT CLASS=“molecule”>H<SUB>2</SUB>0</FONT>
```

Mathematicians will represent this molecule in MathML (MathML, 2001) as:

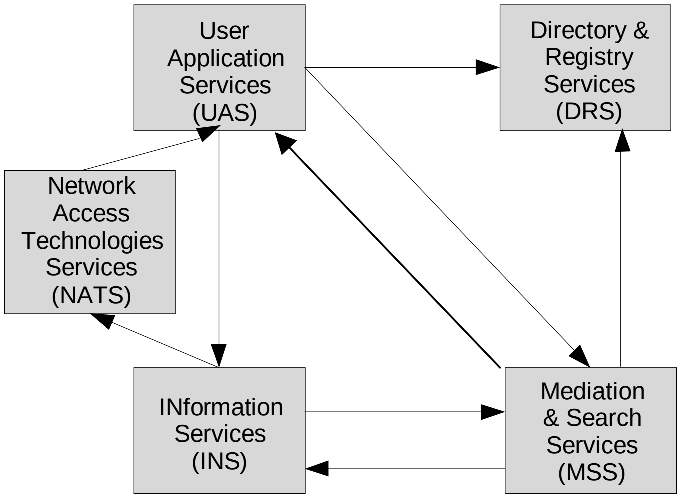

Figure 2: GUARDIANS Functional Architecture

```
<math>
  <msub>
     <mi>H</mi>
     <mn>2</mn>
  </msub>
  <mi>O</mi>
</math>
```

Whereas chemists might use the MoDL (Molecular Dynamics Language) (MoDL, 1999)

```
<DEFINE type='molecule' name='Water' >
<atom type='O' id='o1' position='0 0 0' />
<atom type='H' id='h1' position='-2 1 0.5'/>
<atom type='H' id='h2' position='2 1 0.5'/>
<bond from='h1' to='o1'/>
<bond from='h2' to='o1'/>
</DEFINE>
```

After investigation, the most suitable model for transforming structured metadata was determined to be a canonical data model rather than providing the transformations between all those formats. This model requires that as a new metadata format be added it is only necessary to create a mapping between the new metadata format and the canonical middle language. This is shown in Figure 1.

The Basic Semantic Registry (ISO/IEC 11179 Part 6, 1997) makes use of this canonical language (CL) to translate from one metadata format into one or more other(s). Most significantly, the BSR is used to map user preferences to content metadata representing the information provider's content, regardless of the content metadata format. It is the job of the BSR to provide mappings between the user profile and each of these target metadata types.

As the BSR is responsible for all XML metadata conversions, it uses XSLT (XSL Transformations, 1999). Of course because XSLT is quite a specific language, it is sometimes necessary to re-format the data to enable its successful processing and translation.

4 FEDERATED SEARCH STRATEGY

The personalised resource discovery strategy is performed by a software component set called Mediation and Search Services (MSS, figure 2), which comprises a Mediator and a Search Service. For the purpose of the federated search strategy we will only concentrate on the components of interest for the search: the Mediation and Search Services.

The resource discovery strategy works on a three-phase basis. The initial pass, performed by the Mediator, is to select INSPs who can provide content suitable for the learner initiating the search. The second phase of the search is performed by the Search Service component of the MSS, which searches on the user selected INSPs for the relevant content. The third phase is the results collection and ranking, delivering the pool of results back to the user.

4.1 Finding the relevant INSP

This is achieved through the Mediator component matching the user's profile (GUP) with the INSP profiles stored in the Directory and Registry Service (DRS). The underlying technology used in this matching is XSLT (XSL Transformations, 1999).

The Mediator serves as the middle tier interacting with the INSP profile, the graphical user interface (UAS in figure 2) and the actual Search Service. From the UAS, in interaction with the user profile, the user defines his current preferences and keywords to formulate a search request. Those user preferences are then mapped to the canonical language. A sample extract is as follows:

Table 3: GUP preferences mapping to the canonical language

GUP	CL
//accessibility/language= 'en'	//general/description/ language='en'
//accessibility/preference /refinement='video/quic ktime'	//general/technical/prefer ence='video/quicktime'
//accessibility/preference s/preference/type='Oper ating System'	//general/technical/title= 'Operating System

Those types of mappings are expressed in XSLT (Bradley, 2000) as simple 'select-match' expressions, which are used to select the current element to match it to a particular target element.

Once this mapping is achieved, an XPath query (Clark and DeRose, 1999) can be formulated in the canonical language with the relevant user information. This query formulation process is similar to the second phase's, hence will be described in the next section.

4.2 Finding the relevant resources

For each of the INSPs selected by the Mediator, a subset of the INSP profile is retrieved from the DRS. Through this mechanism the location, type (query type) and format (metadata schema) of each INSP's metadata repository is discovered.

In the next step, preferences are processed using an XSLT document similar to the one defined in Table 3 but this time with the INSP information. Once the previous step is performed a general query is formulated, based on the canonical language (CL). However, this query cannot be executed against any repository since it is schema "unaware". To obtain a schema specific query, the Search Service has to call the BSR. A query is built according to the type and format of the repository. An example of such query for an English-speaking user requesting a search with 'universe' as a keyword and 'quicktime' as his preferred format could be:

```
<search>
 <repositoryname>rep 1</repositoryname>
 <querystring>[!CDATA
  /lom/general/language[text()="en"] |
  /lom/general//lom [(/lom/general/title/
  langstring[contains(text(),"universe") or
  /lom/general/description/langstring[
  contains(text(),"universe")]]
   /lom/technical/format[text()="video/
   quicktime"]
 </querystring>
 <querylanguage>XPATH</querylanguage>
</search>
```

In this example, 'universe' is the only input from the user, language and format are derived from the user's profile. This query is then delivered to the location specified and results added to the overall pool of results.

4.3 Results handling

A result manager component of the MSS is in charge of collecting, ranking and re-formatting the results before they are sent back to the user. The idea is to create a result collection, which is parsed and used for ranking. This result collection is ranked according to the preferences, keywords, etc. This was developed in such a way to allow extension of the rating algorithm, for instance to use other fields from the User Profile. Once accomplished, the results are stored into the following structure:

```
<resourceResult id="r.1">
<title>The Sun</title>
<description>The movie describe some of the
different satellites studying the Sun: The
hubble telescope, Solar max, and the very
large array in new Mexico; and information
about the solar system.</description>
<location>http://www.ist-guardians.org/
astronomy/Sun.mov</location>
<ranking>0.68</ranking>
<format>application/quicktime movie</format>
<inspName>INSP 1</inspName>
</resourceResult>
```

This example represents the XML structure of the results, which are returned by the Search Service. Basically, results are represented the same way as a normal web search engine with the following features describing the XML document above:

- A title to identify the result found.
- A description, to briefly identify what the resource is about.
- A location, to provide a link to the resource found.
- A ranking, as a degree of relevance of this result that will be displayed to the user later on.

- A format representing the resource format for display purposes. For instance, the type of movie, audio, text to know the platform and requirement for displaying it to the user.
- An INSP name, in case the user wants to come back to this particular INSP later on. (Part of the action history in the user profile)

The results discovered are rated using a simplistic mechanism as follows:
1: an arbitrary base score (30%) is given for all results (since it has been found, matching a minimum of one keyword).
2: a score is given for each keyword found in the title (20%).
3: a score is given for each keyword found in the description (6%).
4: a bonus score is given if the title or description contains more than one keyword (10 and 5%).
5: a score is given for the preference level at which the result was found (3%).

A greater number of preferences matched, means a better, more personalised result for the user.

5 CONCLUSION AND FUTURE WORK

As the number of resources on the Internet continues to grow, user faces the issue of information overload. In this situation, reliable search tools that enable personalisation of resource discovery and information retrieval becomes crucial. This is specifically true in the E-Learning domain, where users (students or researchers) have a clear need to retrieve information relevant to their preferences (language, interest, platforms, etc). Simple text-based search engines are not sufficient anymore. Several improvements have been made in the past few years showing a better search accuracy. For instance, Google (Page, 1998) now provides a new improved HTML based search engine with regard to quicker delivery and better search accuracy. Although such tools are successful, it is only a partial solution, it only delays the problem as more and more information (educational, recreational, etc) becomes available and as the number users still increases.

A complete solution to this problem is the combination of current and emerging technologies and standards to offer better interaction between the user and federated search strategies. This solution, presented in this paper, has been researched and implemented by the author in the E-Learning context of the GUARDIANS European project as part of its Mediation Service component. It offers a very efficient search facility that enhances the user's learning experience through the usage of smart functionalities (components and resources) for discovery and retrieval of educational resources. Although this work is based on E-Learning resources, one could easily fit it into the overall context of information retrieval on the Internet.

Although the search strategy using XPath performed successfully, it was felt that the queries were limited with regards to performance and flexibility. Indeed, to incorporate most of the user preferences, the XPath was complicated and major results reformatting had to be performed. For this reason this initial work was extended and an XML querying solution using XQuery (Boag et al., 2002) was proposed in (Rousseau, Leray, 2003).

Another important point to consider is that the matching between user profile and metadata was limited but it would be possible to include more information about the user. For example, the search could use the affiliation to only search for INSPs with which the user is affiliated. It could also use qualification details to only retrieve information relevant to the user's previous qualification, diplomas and other licenses or certifications. However, the user profile might still be extended with new features according to future development in the E-Learning community.

The ranking algorithm implementation was an interesting part of the work. However, with time constraints and the lack of information about how to actually develop a precise ranking mechanism. It was first developed as a basic algorithm, only taking care of the keywords. For a possible commercialisation purpose, it should have taken into account a few more feature from the user profile, specifically provide a complex algorithm using weighting attributes. It is expected that the authors will be involved into future research in the area of stopwords, tf.ifd, coordination level, etc (Baeza-Yates, 1999).

Nowadays, there is a lot of hype around agents. In the context of this work, an agent would enable automatic discovery of the user's platform capabilities. This issue was encountered while populating the user profile with some device information. Hence, a typical user would be unaware of its device(s) capabilities and will require an automatic detection mechanism. Also, a good balance between static and dynamic profile information should be researched further to present the user with a reliable profile that requires the minimum efforts to fill in. A form of interaction between the profiling mechanism and the user is therefore crucial to get some feedbacks to the system.

The abstract representation of data on the World Wide Web, also known as the Semantic Web

(Berners et al., 2001), is gradually taking over the current HTML based Web. In the near future, the development in the area will lead in significant new possibilities to enhance and process the data that they merely display at present. Hence, programs will be able to exchange information with others and automate services. Such functionalities together with the flexibility of XML will be promoted using agents, that will help users tracking down resources and match them against their criteria and preferences.

REFERENCES

Baeza-Yates, R. and Ribeiro-Neto, B. Modern Information Retrieval, *Addison-Wesley, ACM Press, 1st edition (May 1999*). ISBN: 020139829X

Berners-Lee T., Hendler J. and Lassila O. The Semantic Web. A new form of Web content that is meaningful to computers will unleash a revolution of new possibilities. *Scientific American, May 17, 2001*.

Boag, S., Chamberlin, D., Fernandez, M.F., Florescu, D., Robie, J., Simeon, J., and Stefanescu, M., eds 2002. XQuery 1.0: An XML Query Language. *W3C Working Draft,* http://www.w3.org/TR/2002/WD-xquery-20020430/

Bradley, N., September 7, 2000. The XSL Companion. *Addison-Wesley Pub Co; 1st edition*. ISBN: 0201674874

Chang, B., Scardina, M., Karun, K., Kiritzov, S., Macky, I., Novoselsky, A. and Ramakrishnan, N., 2000. ORACLE XML Handbook. *Oracle Press.* ASIN: 007212489X

Clark, J., and DeRose, S. eds 1999. XML Path Language (XPath) version 1.0. *W3C Recommendation.* http://www.w3.org/TR/xpath

Day, M., Summer 2001. Metadata in a nutshell. *Article published in Information Europe, p. 11. Information Europe quarterly magazine of EBLIDA.*

Florescu, D., and Kossmann, D., 1999. A Performance Evaluation of Alternative Mapping Schemes for Storing XML Data in a Relational Database. *Technical report no. 3680, INRIA*, Le Chesnay Cedex, France.

GESTALT, Getting Educational Systems Talking Across Leading-Edge Technologies, EU project (started in October 1998) Web Site: http://www.fdgroup.com/gestalt/

GUARDIANS project, 2000. *Gateway for User Access to Remote Distributed Information and Network Services, started in October 2000*. Web site: http://www.ist-guardians.tv/

Howes, T., Yeong, W. and Kille, S, 1995. Lightweight Directory Access Protocol (LDAP). *The Internet Society, RFC 1770*.

IEEE Learning Technology Standards Committee WG12: Learning Object Model, 6 March 2002. http://ltsc.ieee.org/wg12/

IEEE PAPI, 01 February 2002. *IEEE Public and Private Information draft 8 specifications*. Available at: http://www.edutool.com/papi/

IMS LIP version 1.00, 18 March 2001. *Learner Information Package Specification*. Available at: http://www.imsglobal.org/profiles/index.cfm

IMS Learning Resource Metadata Specification, 01 October 2001. http://www.imsglobal.org/metadata/index.cfm

ISO/IEC 11179 Part 6, 1997. *Information technology - Specification and standardization of data elements - Part 6: Registration of data elements.* Available online at: http://www.iso.ch/

Kanne, C.-C., and Moerkotte, G., 2000. Efficient Storage of XML Data. *Proceeding of the International Conference on Data Engineering*. San Diego, California, p. 198.

MathML, W3C math specification, 21 February 2001, http://www.w3.org/Math

MoDL - Molecular Dynamics Language project, 29 July 1999, http://xml.coverpages.org/modl.html

Page, L., Brin, S. The anatomy of a large-scale hypertextual Web search engine, *Computer Networks and ISDN Systems, Volume 30, Issues 1-7, April 1998, Pages 107-117.*

RDS. Distributed Systems Technology Centre: Research Data Unit. *Resource Discovery – A Definition*. Available from http://www.dstc.edu.au/RDU/rd_define.html

Rousseau, B., Browne, P., Malone, P., ÓFoghlú, M., 2003. User Profiling for Content Personalisation in Information Retrieval. *In the ACM Symposium on Applied Computing, special track on Information Access and Retrieval*. ACM Press

Rousseau, B., Leray, E., ÓFoghlú, M., 2003. Metadata & Information Management issues in XML-based Mediation. *In ICEIS'03 Fifth International Conference on Enterprise Information Systems, Angers, FRANCE, pages 375-379*. ICEIS Press.

RSLP. Research Support Libraries Programme British initiative, September 1999, http://www.rslp.ac.uk/

Spark, K. Jones, Artificial Intelligence: What can it offer to Information Retrieval. *In Proceedings of the Informatics 3, Aslib, ed., London, 1978.*

Suryanarayana, L., Hjelm, J., May 7-11, 2002. Profiles for Situated Web. Johan. *In WWW 2002, Honolulu, Hawaii, USA*. ACM Press 1-58113-449-5/02/0005.

TV Anytime forum, formed in 1999. Web site, available at: http://www.tv-anytime.org/

UKOLN, centre of expertise in digital information management, University of Bath, England. http://www.ukoln.ac.uk/

vCard, December 1996. The Internet Mail Consortium (IMC), available at: http://www.imc.org/pdi/vcardoverview.html

XML, the eXtensible Markup Language 1.0 second Edition. *W3C Recommendation 6 October 2000*, http://www.w3.org/TR/REC-xml

XSL Transformations (XSLT), version 1, *W3C Recommendations, 16 November 1999*, http://www.w3.org/TR/xslt