

Adaptive Personalisation in Self e-Learning Networks

Kevin Keenoy, Alexandra Poulouvassilis, George Papamarkos, Peter T. Wood
School of Computer Science and Information Systems, Birkbeck, University of London
{*kevin, ap, gpapa05, ptw*}@*dcs.bbk.ac.uk*

Vassilis Christophides, Aimilia Magkanaraki, Miltos Stratakis
Institute of Computer Science, Foundation for Research and Technology – Hellas (FORTH-ICS)
{*christop,aimilia,mstratak*}@*ics.forth.gr*

Philippe Rigaux, Nicolas Spyratos
Laboratoire de Recherche en Informatique, Universite Paris-Sud
{*rigaux,spyratos*}@*lri.fr*

Abstract

This paper presents some of the personalisation services designed for *self e-learning networks* in the SeLeNe project. A self e-learning network consists of web-based learning objects that have been made available to the network by its users, along with metadata descriptions of these learning objects and of the network's users. The architecture of the network is distributed and service-oriented. The personalisation facilities include: querying learning object descriptions to return results tailored towards users' individual goals and preferences; the ability to define views over the learning object metadata; facilities for defining new composite learning objects; and facilities for subscribing to personalised event and change notification services.

Keywords: Self e-Learning Networks, Personalisation, RDF View definition, Learning trails, RDF ECA Rules

1. BACKGROUND - THE SELENE PROJECT

The increasing use of global telecommunications and the internet for learning has led to the formation of diverse communities of learners that are geographically distributed and whose members have heterogeneous educational backgrounds. Now that such virtual learning communities (VLCs) can also be built upon Grid technologies [19] there is an increasingly urgent need for the development of support tools for such communities.

At the same time, the number of learning resources available on the web is continuously increasing, indicating the web's enormous potential as a significant resource of educational material both for learners and instructors.

The SeLeNe (*self e-learning networks*) project¹ has investigated the design of tools for supporting learning communities, matching learners' needs with the educational resources potentially available on the web. SeLeNe relies on semantic metadata describing *learning objects* (LOs), and provides services for the discovery, sharing and collaborative creation of LOs that facilitate syndicated and personalised access to such resources. The LO metadata descriptions and the schemas that these conform to form SeLeNe's *LO information space*. Users need to be able to query the LO information space in order to locate resources appropriate for their specific learning or teaching needs, and also need to be able to define personalised views over this potentially large number of heterogeneous resources.

We envisage SeLeNes being created by and serving different communities of users, each SeLeNe supporting a specific community. These communities could occur, for example, in an educational institution, as a community of professional practice or in a geographical region. In the most general case learners organise themselves into online VLCs according to their own criteria such as needs and interests.

¹<http://www.dcs.bbk.ac.uk/selene>.

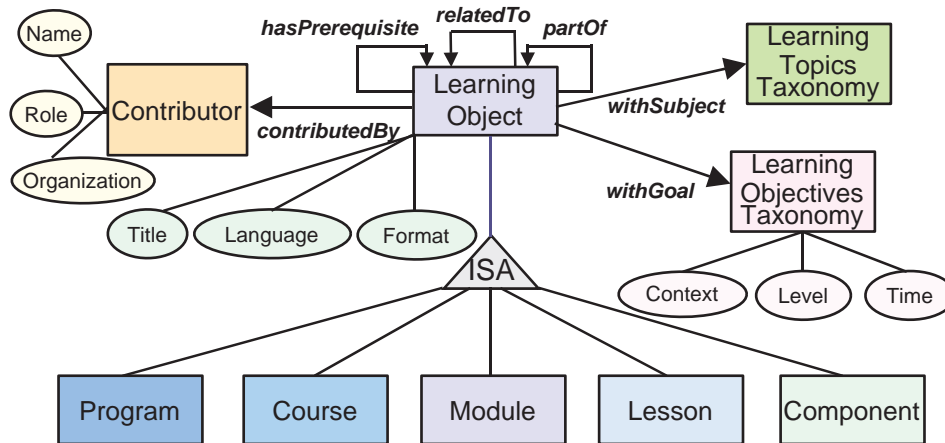


FIGURE 1: Overview of the LO Descriptive Schemas in SeLeNe

The LOs described by a SeLeNe are electronic, web-based, sharable chunks of reusable learning content that have been explicitly made available to the network by its users via metadata describing the LOs (including their URIs). We call the process of metadata submission *registering a LO with a SeLeNe*. The LOs made available to the network are not necessarily hosted by the individual members of the SeLeNe (although they may be) – they can reside anywhere reachable on the network infrastructure being used (web, Grid, etc.). A crucial point is that a SeLeNe does not manage the LOs themselves, but instead facilitates access to them by managing their metadata descriptions.

In order to enable effective search for LOs in a SeLeNe, LO descriptions conform to e-learning standards such as IEEE/LOM (Learning Object Metadata) [8, 16], and also employ topic-specific taxonomies of scientific domains such as ACM/CCS (Computing Classification System) [2] or taxonomies of detailed learning objectives. All LO schemas and descriptions are represented in the Resource Description Framework (RDF) or Resource Description Framework Schema (RDFS) languages [11, 3], which offer advanced modelling primitives for the SeLeNe information space. (In this paper we use ‘RDF/S’ to mean ‘RDF and/or RDFS’.)

Figure 1 illustrates the main concepts and properties of the RDFS schemas employed by SeLeNe to capture the semantics of existing e-learning standards. The information content of a LO can be described using attributes such as *title*, *language*, *format*, etc., or one or more terms from a topic-specific taxonomy like ACM/CCS (for example, a LO about C++ could have a *withSubject* property referencing the ACM/CCS classification *D.3.2.11*, which is the taxonomic path representing *Software.Programming.Languages.Language.Classifications.ObjectOriented.Languages*).

The schema also allows description of the granularity of a LO (e.g., *Course*, *Lesson*), its relationships with other LOs (e.g., *hasPrerequisite*, *partOf*), and its relationships with other classes of resource (e.g., *contributedBy*). A taxonomy such as Bloom’s *Taxonomy of Educational Objectives*, combined with a topics taxonomy, can be used to describe the desired learning outcomes of a LO. For example, a LO *withGoal Application.Use* (from Bloom’s taxonomy) and *withSubject D.3.2.11* (from ACM/CCS) is thereby described as having the educational objective “to be able to use an Object-Oriented language”.

No single architectural design will be suitable to support the diverse communities who may wish to use a SeLeNe. The needs of any particular learning community are addressed by allowing them to choose the most suitable deployment of the service-based architecture that we have defined for their network infrastructure. The services can be deployed in a centralised, mediation-based or peer-to-peer fashion.

In the most general deployment scenario, a SeLeNe network will consist of a number of peers, each of which will support some subset of the full set of SeLeNe services. Figure 2 illustrates the general architecture of a SeLeNe. As long as there is at least one instance of each of the services available to the network then a SeLeNe can be formed. The facilities that are the focus of this paper are provided by the User Registration,




FIGURE 2: SeLeNe Service Architecture

LO Registration, Trails and Adaptation, Event and Change Notification, and View services. We refer the reader to [20] for further discussion of the service architecture and other services.

The users of a SeLeNe will include *instructors*, *learners* and *providers* of LOs – a single person could play each of these roles at different times. New users will be able to join a SeLeNe by contacting any peer that provides the User Registration service. When registering, users will supply information about themselves and their educational objectives in using this SeLeNe. This information is stored in their *personal profile*, which is an RDF description conforming to a number of RDFS schemas.

LO authors maintain control of the content they create and will be free to use any tools they wish to create their LOs before registering them. We call such LOs created externally to SeLeNe, *atomic LOs*. Users will also be able to register new *composite LOs* that have been created as assemblies of LOs already registered with the SeLeNe. The taxonomical descriptions for these LOs can be automatically derived by the SeLeNe from the taxonomical descriptions of their constituent LOs and this process is discussed in Sect. 2.

The SeLeNe will provide facilities for defining personalised *views* over the LO information space, as discussed in Sect. 3. The SeLeNe will also provide browsing and searching facilities over the LO information space, which return results tailored towards users' individual goals and preferences. This personalisation is based on an adaptive user profile, which includes a history of recent user activity that can be mined for information to help adaptation to the user. These facilities are discussed in Sect. 4.

A user may wish to be notified of changes made to a LO's description. SeLeNe will provide such automatic change detection and notification facilities by comparing the new and old description of the LO. SeLeNe will also support personalised notification services depending on users' profiles and will automatically update users' profiles to reflect recent interactions with the SeLeNe. Provision of these facilities is discussed in Sect. 5.

2. REGISTRATION OF LOS

Registration of a LO with a SeLeNe consists of providing a metadata description including the URI of the LO. An important feature of SeLeNe's LO Registration service will be the ability to automatically infer the *taxonomical* part of the description for a composite LO o , from the taxonomical descriptions of its component LOs o_1, \dots, o_n . This description, which 'summarises' the taxonomical descriptions of the parts of o , is called the *implied taxonomical description* (ITD) of o . This is used to augment the *publisher taxonomical description* (PTD) – the set of terms supplied by the LO's provider (which can be empty) – to derive the final taxonomical description used to register the composite LO.

The ITD of a composite LO o expresses what its parts have in common – only terms reflecting the content of all parts are included in the ITD. Inclusion of an LO o_1 as part of a composite LO o may thus lead to the generation of different ITDs for o depending on what the 'companion' parts to o_1 are within o .

The ITD of a composite LO o composed of parts o_1, \dots, o_n with descriptions D_1, \dots, D_n is computed by a simple algorithm that takes the cartesian product of D_1, \dots, D_n , computes the least upper bound of each n -tuple and then 'reduces' the resulting set of terms by removing all but the minimal terms according to the subsumption relation \preceq between the terms of the taxonomy. The overall taxonomical description of a LO o is computed by another simple algorithm: if o is atomic then its taxonomical description is just its PTD. Otherwise its taxonomical description is recursively computed from its PTD and the taxonomical descriptions of its constituent parts. Readers are referred to [18] for more details of both algorithms.

3. DECLARATIVE QUERIES AND VIEWS

Finding LOs in a SeLeNe will rely on RQL [9], a declarative query language offering browsing and querying facilities over the RDF/S descriptions that form the SeLeNe information space.

As well as the advanced querying facilities provided by an expressive RDF/S query language such as RQL, personalisation of LO descriptions and schemas is also needed. For instance, a learner might want LOs presented according to his/her educational level and current course of study. To enhance the SeLeNe user's experience we need the ability to personalise the way the information space can be viewed, by providing simpler virtual schemas that reflect an instructor's or learner's perception of the domain of interest. This can be done by using the RVL language [14], which provides this ability by offering techniques for the reconciliation and integration of heterogeneous metadata describing LOs, and for the definition of personalised views over a SeLeNe information space.

One of the most significant features of RVL is its ability to create virtual schemas by simply populating the two core RDF/S metaclasses *rdfs:Class* and *rdf:Property*. A SeLeNe user can then easily formulate RQL queries *on the view* itself. RVL can also be used to implement advanced user aids such as *personalised navigation* and *knowledge maps*. View definition in SeLeNe is discussed in detail in [13].

4. TRAILS AND QUERY ADAPTATION

Personalisation of query results will rely on the *personal profile*. This is an RDF description of the user including some elements from existing profile schemas (PAPI-Learner and IMS-LIP), some of our own elements (we have defined RDFS schemas for expressing competencies, learning goals and learning styles), and a history of user activity that will allow the profile to adapt over time, automatically updating the information therein by means of a generic set of Event-Condition-Action (ECA) rules (see Sect. 5) associated with all SeLeNe profiles. The learner also has a *messages* property with a *Notifications* class as its target. This is to store personal notifications (of new users and new or updated LOs) for the user. Figure 3 shows a simplified version of SeLeNe's personal profile schema.

Although the underlying query mechanism in SeLeNe is RQL, users will mostly search for LOs using simple keyword-based queries. Search results will be personalised by *filtering* and *ranking* the LOs returned according to the information contained in the user's personal profile. This is done by the Trails and Adaptation service, which constructs personalised RQL queries for execution and generates personalised rankings of query results by matching the personal profile against relevant parts of the LO descriptions.

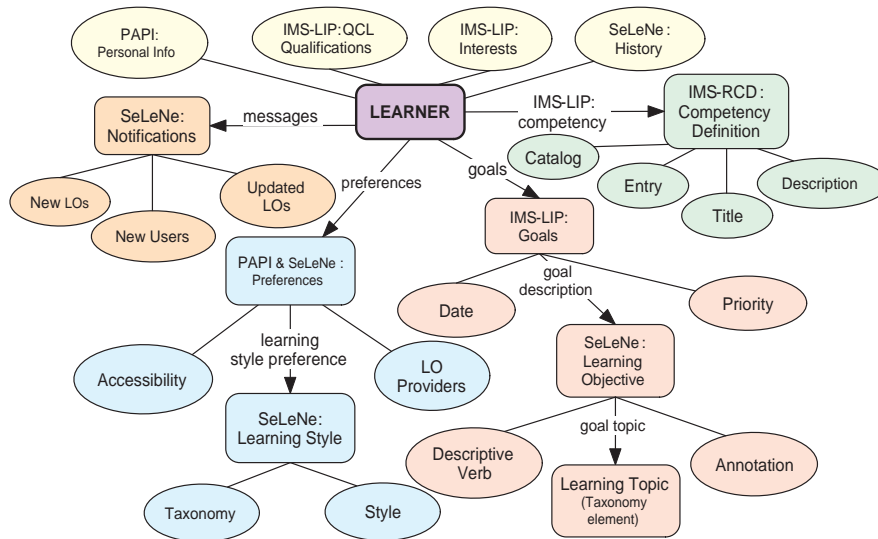


FIGURE 3: SeLeNe's Personal Profile Schema

Initial filtering takes place before a query is evaluated: all keyword-based queries submitted to the SeLeNe will be routed through the Trails and Adaptation service, which will construct corresponding personalised queries. The user's original query will then be re-formulated and/or annotated to reflect elements of their personal profile. For example, if the profile records that the user only speaks English, all searches they enter can be augmented with the annotation *lang:en*, ensuring that all returned LOs will be in English. The augmented keyword query will then be translated into an appropriate RQL query and executed by the Query service. The reader is referred to [10] for the full SeLeNe profile schema and for discussion of the translation of keyword-based queries into structured queries in SeLeNe.

The next step is to rank the set of LO descriptions in order of relevance to the user. Relevance will be judged against the combination of the personal profile and the original query. The information contained in the SeLeNe personal profile allows the following factors to be taken into consideration when judging relevance: relevance of the LO to the query; how well the LO caters for the user's accessibility requirements; whether the user has the prerequisite knowledge and experience to be able to tackle the LO; how well the learning objectives of the LO match the user's learning goals; if the user's learning styles are those catered for by the LO; if the user is likely to prefer it for other reasons (e.g., it is by a preferred author); and the user's most recent activity (as contained in the history section of their profile).

The different sections of the personal profile can be matched in a focussed way against relevant sections of the LO descriptions. For one LO to be more relevant to the user than another it must meet more of the user's different preferences, or match preferences to a greater degree. For example, if the user has OOL as one of their *goal.topic*'s then a LO including OOL in its taxonomical description will score well on relevance to the user. If the LO description additionally includes the information that it caters for one of the user's learning styles then it will be considered a better match again. The best algorithm and weightings to use for this ranking needs to be determined empirically, and may well need to be adaptive.

Once a personalised ranking of the remaining LOs has been generated the search results will then be returned to the user. SeLeNe will give the user the option of having their query results presented as a list of *trails* of LOs, where a trail is a suggested sequence of interaction with the LOs. These trails will be automatically derived from information contained in the LO descriptions about the semantic relationships between LOs. In the SeLeNe project we defined an RDF representation of trails whereby they are defined as a sub-class of the RDF *Sequence* (a sequence of LOs) with two associated properties, *name* and *annotation*, that provide additional information about the pedagogic use of the trail. In subsequent work, as part of the TRAILS project², we have developed a richer metadata schema for learning trails that expands the

²<http://www.dcs.bbk.ac.uk/trails>

enumeration of possible types of trail and additionally includes (optional) timing information, topic, level, and other information enabling more complex learning design [21].

The history part of a user profile is also a trail. It is key to the dynamicity of the profile, and thus to SeLeNe's adaptivity to individual users. As users interact with LOs via a SeLeNe the history of their accesses to LOs is recorded as a trail. The SeLeNe can use information stored in the semantic metadata about the visited LOs to infer current user interests, learning goals and learning styles and then automatically update the user profile to reflect these inferences. In this way the user is spared from having to constantly update their profile by manually entering information – the system automatically adapts to users' preferences simply by monitoring their behaviour. A user profile is updated and maintained by a set of ECA rules that are created whenever a new user profile is registered with the system. Such rules are the subject of the following section.

5. EVENT AND CHANGE NOTIFICATION

SeLeNe's reactive functionality includes features such as:

- propagating changes in the taxonomical description of a LO to those of any composite LOs depending on it;
- propagating changes in a learner's history of accesses of LOs to the learner's personal profile;
- notifying users of the registration of new LOs of interest to them;
- notifying users of changes in the description of resources of interest to them.

This reactive functionality is provided by means of ECA rules over SeLeNe's RDF/S metadata. These ECA rules will be automatically generated by the higher-level Presentation and Application services of the SeLeNe architecture. An ECA rule generated at one site of the network might be triggered, evaluated, and executed at different sites. Within the event, condition and action parts of ECA rules there might or might not be references to specific RDF resources, i.e. ECA rules may be resource-specific or generic.

We refer the reader to [17] for a detailed description of the syntax and semantics of *RDFTL*, our language for specifying ECA rules on RDF. Here we give two examples of RDFTL rules, which refer to the LO Schema illustrated in Figure 1 and the Personal profile schema illustrated in Figure 3.

Firstly, if a LO is inserted whose subject is the same as one of user 128's goal topics, then the following rule adds a new arc linking the newly inserted LO into the `new_LOs` collection in user 128's personal messages:

```
ON INSERT resource()
  AS INSTANCE OF LearningObject
  IF $delta/target(lom:subject) = resource(http://www.dcs.bbk.ac.uk/users/128)
    /target(ims-lip:goal)
    /target(ims-lip:goaldescription)
    /target(selene:goaltopic)
  DO LET $new_los := resource(http://www.dcs.bbk.ac.uk/users/128)
    /target(selene:messages)
    /target(selene:new_LOs)
  IN INSERT ($new_los, seq++, $delta);;
```

Here, the event part checks if a new resource belonging to the `Learning Object` class has been inserted. The condition part checks if the inserted LO has a subject which is the same as of one user 128's goal topics. The `LET` clause in the rule's action defines the variable `$new_los` to be user 128's new LOs collection. Finally, the `INSERT` clause inserts a new arc from `$new_los` to the newly inserted LO (we use the syntax `seq++` to indicate an increment in the collection's element count).

As a second example, if a change is made to any property of a LO whose subject is the same as one of user 128's goal topics, then the following rule adds a new arc linking user 128's `updated_LOs` collection to the modified LO:

```
ON UPDATE (resource() AS INSTANCE OF LearningObject, _, _->_)
  IF $delta/target(lom:subject) = resource(http://www.dcs.bbk.ac.uk/users/128)
    /target(ims-lip:goal)
    /target(ims-lip:goaldescription)
```

```

    /target(selene:goaltopic)
DO LET $updated_los := resource(http://www.dcs.bbk.ac.uk/users/128)
    /target(selene:messages)
    /target(selene:updated_LOs)
IN INSERT ($updated_los, seq++, $delta);;

```

Those peers in a SeLeNe that support the Event and Change Notification service will run an *ECA Engine*. Such ‘superpeer’ servers may coordinate a group of further ‘peer’ servers (as well as themselves playing the role of a ‘peer’). Each peer may host a fragment of the global RDF/S descriptions in the SeLeNe network, stored in a local RDF/S repository (in our current implementation the repository is RDFSuite [1]). The fragment of the global RDFS schema stored at a peer may change as a result of changes in the peers’ RDF/S descriptions. Peers notify their coordinating superpeer of any updates to their local RDF/S repository, so that the superpeer is able to undertake Event Detection (see below).

The ECA engine at a superpeer consists of several subcomponents:

- The *RDFTL Language Interpreter* takes as input an RDFTL rule definition and registers it in a Rule Base. The Language Interpreter consists of a *Parser* which verifies the syntactic correctness of the rule, and a *Translator* which translates any path queries within the rule into the query syntax of the underlying RDF repository. In our present implementation, RDFTL path expressions are translated into RQL.

Whenever a new ECA rule, r , is registered at a peer, P , it will be sent to P ’s superpeer for registration and storage. From there, r will be forwarded to all other superpeers (using the routing service – see below) and a replica of it will be stored at those superpeers that are relevant to it, i.e. where an event may occur in a local RDF/S repository that will trigger r ’s event part.

- The *Event Detector* detects the occurrence of events within a local RDF/S repository and the triggering of rules registered within the local rule base. The Event Detector determines which rules have been triggered by an update to a local repository by invoking that repository’s Query Service to evaluate the event queries of rules that may have been triggered.
- The *Condition Evaluator* determines which of the triggered rules should fire. It does this by submitting appropriate queries to the superpeer’s Query Service to determine which of the triggered rules’ conditions are true. This may require distributed query processing across a number of peers and superpeers, which is coordinated by the superpeer’s Query service.
- If a triggered rule’s condition evaluates to true, the *Action Scheduler* generates from the action part of the rule a set of updates to be sent for execution to the appropriate local peers and other superpeers.
- The *Routing Service* maintains information about the superpeer’s immediate neighbours in the network and the message transmission pathways between them.

6. COMPARISON WITH RELATED SYSTEMS

In this section we briefly compare the functionality of SeLeNe with that of four related systems: UNIVERSAL³, Edutella⁴, Elena⁵, and SWAP⁶.

UNIVERSAL [7] is a business-to-business brokerage platform aiming to support higher education institutions in the exchange of learning resources. It allows institutions to advertise their learning resources, and provides an RDF/S-based catalog which can be browsed to find and access learning resources. Edutella [15], which is based on the JXTA peer-to-peer framework [6], provides a peer-to-peer infrastructure for connecting peers supporting different types of repositories, query languages, and metadata schemas. Each peer implements a number of basic services such as querying, replication and mapping between different schemas. Elena [22] provides a mediation infrastructure for learning services, where these might be provided by assessment tools, Learning Management Systems, etc. It includes dynamic learner profiling using ‘personal learning assistants’. SWAP [4, 5] does not address e-learning specifically but is investigating the integration of semantic web and peer-to-peer technologies in order to support knowledge sharing. It is

³www.ist-universal.org

⁴edutella.jxta.org

⁵www.elena-project.org

⁶swap.semanticweb.org

developing technology both for allowing users individual views of knowledge and for effective sharing of knowledge.

The novel aspects of SeLeNe compared with these systems include:

- collaborative creation and semi-automatic description of composite LOs, which do not seem to be addressed specifically by any other system;
- declarative views over combined RDF/RDFS descriptions i.e. over both the LO descriptions and their schemas (SWAP provides users with the ability to define views over RDF descriptions only, while Edutella and Elena do not ensure the compositionality of queries with views and mappings);
- personalised event and change notification services (event notification services are used in Edutella [12] to assist in rule-based clustering of peers, where the events are caused by peers connecting to, or disconnecting from, a super-peer; SeLeNe's events, on the other hand, operate at the level of LO or user descriptions);
- use of trails recording histories of LO accesses to automatically update user profiles.

7. CONCLUDING REMARKS

This paper has described several novel techniques for providing the personalisation services of *self e-learning networks*. We are currently finishing the implementation and integration of some of these services, and we then plan to deploy and evaluate them within a Grid/peer-to-peer architectural framework. In particular, our current implementation of Event and Change Notification services exploits the JXTA peer-to-peer framework [6] in order to deliver active functionality in a distributed peer-to-peer environment over RDF e-learning metadata repositories.

A number of open issues remain. Firstly, there is as yet no standard query or update language for RDF, although we believe that the RQL, RVL and RDF ECA languages we have described here provide sound and expressive foundations for the development of such standards. Another important open area is combining ECA rules with transactions and consistency maintenance in RDF repositories, particularly in a peer-to-peer environment. Thirdly, our algorithms for personalised ranking of query results need to be empirically evaluated. Finally, the design of user interfaces enabling easy and intuitive access to SeLeNe's advanced personalisation services will be crucial – users will need to be shielded from the complexities of RDF and the RQL, RVL and ECA languages, and also from the complex taxonomies of topics, competencies and goals in use by the system.

REFERENCES

- [1] S. Alexaki, V. Christophides, G. Karvounarakis, D. Plexousakis, and K. Tolle. The ICS-FORTH RDFSuite: Managing Voluminous RDF Description Bases. In *Proc. of the 2nd International Workshop on the Semantic Web (SemWeb 2001)*, 2001.
- [2] Association for Computing Machinery. The ACM computing classification system (1998 version). See <http://www.acm.org/class/1998/>, 1998.
- [3] D. Brickley and R. Guha. RDF vocabulary description language 1.0: RDF schema. See <http://www.w3.org/TR/rdf-schema/>, 2004.
- [4] J. Broekstra, M. Ehrig, P. Haase, F. van Harmelen, A. Kampman, M. Sabou, R. Siebes, S. Staab, H. Stuckenschmidt, and C. Tempich. A Metadata Model for Semantics-Based Peer-to-Peer Systems. In *Proc. of the 1st Workshop on Semantics in Peer-to-Peer and Grid Computing, at WWW2003*, 2003.
- [5] M. Ehrig, C. Tempich, J. Broekstra, F. van Harmelen, M. Sabou, R. Siebes, S. Staab, and H. Stuckenschmidt. SWAP: Ontology-based Knowledge Management with Peer-to-Peer Technology. In *Proc. of the 4th European Workshop on Image Analysis for Multimedia Interactive Services (WIAMIS)*, 2003.
- [6] L. Gong. JXTA: A network programming environment. *IEEE Internet Computing*, 5:88–95, 2001.
- [7] S. Guth, G. Neumann, and B. Simon. UNIVERSAL: Design Spaces of Learning Media. In *Proc. of the 34th Hawaii International Conference on System Sciences (HICS)*, Maui, Hawaii (USA), January 2001.
- [8] IEEE-LTSC. Draft standard for Learning Object Metadata. See <http://ltsc.ieee.org/wg12/files/LOM.1484-12.1.v1.Final.Draft.pdf>, 2002.

- [9] G. Karvounarakis, S. Alexaki, V. Christophides, D. Plexousakis, and M. Scholl. RQL: A declarative query language for RDF. In *Proc. of the 11th International World Wide Web Conference (WWW2002)*, pages 592–603, Honolulu, Hawaii, USA, 2002.
- [10] K. Keenoy, M. Levene, and D. Peterson. Personalisation and trails in Self e-Learning Networks. See <http://www.dcs.bbk.ac.uk/selene/reports/Del4.2-1.4.pdf>, December 2003.
- [11] O. Lassila and R. Swick. Resource Description Framework (RDF) model and syntax specification. See <http://www.w3.org/TR/PR-rdf-syntax/>, 1999.
- [12] A. Löser, W. Siberski, M. Wolpers, and W. Nejdl. Information Integration in Schema-Based Peer-To-Peer Networks. In *Proc. of the 15th Conference on Advanced Information Systems Engineering (CAiSE'03)*, 2003.
- [13] A. Magkanaraki, V. Christophides, G. Karvounarakis, and D. Plexousakis. Views and Structured Querying in Self e-Learning Networks. See <http://www.dcs.bbk.ac.uk/selene/reports/Del43.pdf>, October 2003.
- [14] A. Magkanaraki, V. Tannen, V. Christophides, and D. Plexousakis. Viewing the semantic web through RVL lenses. In *Proc. of the 2nd International Semantic Web Conference (ISWC'03)*, Sanibel Island, Florida, USA, October 2003.
- [15] W. Nejdl, B. Wolf, C. Qu, S. Decker, M. Sintek, A. Naeve, M. Nilsson, M. Palmér, and T. Risch. EDUTELLA: A P2P Networking Infrastructure Based on RDF. In *Proc. of the 11th International World Wide Web Conference (WWW2002)*, 2002.
- [16] M. Nilsson. IEEE Learning Object Metadata RDF binding. See <http://kmr.nada.kth.se/el/ims/metadata.html>, 2003.
- [17] George Papamarkos, Alexandra Poulouvassilis, and Peter T. Wood. RDFTL: An Event-Condition-Action language for RDF. In *Proc. of the 3rd International Workshop on Web Dynamics*, 2004. See www.dcs.bbk.ac.uk/webDyn3/.
- [18] P. Rigaux and N. Spyrtos. Generation and syndication of learning object metadata. See <http://www.dcs.bbk.ac.uk/selene/reports/Del4.1-2.2.pdf>, January 2004.
- [19] N. Romano. Virtual learning communities. *Learning GRID*, (3):3–7, January 2005.
- [20] G. Samaras, K. Karenos, and E. Christodoulou. A Grid service framework for Self e-Learning Networks. See <http://www.dcs.bbk.ac.uk/selene/reports/Del3.pdf>, August 2003.
- [21] J. Schoonenboom, A. Lejeune, J-P. David, M. Turcsányi-Szabó, P. Kaszás, L. Montandon, A. Jones, K. Keenoy, and M. Levene. Visualising trails as a means for fostering reflection. See <http://www.dcs.bbk.ac.uk/trails/docs/D22-02-02-F.pdf>, December 2004.
- [22] B. Simon, Z. Miklós, W. Nejdl, M. Sintek, and J. Salvachua. Smart space for learning: A mediation infrastructure for learning services. In *Proc. of the 12th International World Wide Web Conference (WWW2003)*, 2003.