Contributions to a public e-learning platform: infrastructure; architecture; frameworks; tools

Ambjörn Naeve*, Mikael Nilsson, Matthias Palmér and Fredrik Paulsson

Knowledge Management Research (KMR) group,
Centre for User Oriented Information Technology Design (CID),
Department of Numerical Analysis and Computer Science (NADA),
Royal Institute of Technology (KTH),
Stockholm 100 44, Sweden,
E-mail: amb@nada.kth.se E-mail: mini@nada.kth.se
E-mail: matthias@nada.kth.se E-mail: frepa@nada.kth.se
Website: http://kmr.nada.kth.se
*Corresponding author

Abstract: Life long, flexible, collaborative, and personalized learning are words that are being increasingly used, whenever education is discussed and designed. They express new and important demands on learning architectures, both with regard to pedagogy, organization and technology. Traditional learning architectures are based on teacher-centric and curriculum-oriented "knowledge-push". In this paper we present an infrastructure, an architecture and a number of frameworks and tools that support learner-centric and interest-oriented "knowledge-pull". We see them as a contribution to a Public e-Learning Platform, which can achieve true interoperability based on open source and open international ICT standards

Keywords: e-learning; semantic web; conceptual web; electronic portfolios; peer-to-peer; interoperability.

Reference to this paper should be made as follows: Naeve, A., Nilsson, M., Palmér, M. and Paulsson, F. (2005) 'Contributions to a public e-learning platform: infrastructure; architecture; frameworks; tools', *Int. J. Learning Technology*, Vol. 1, No. 3, pp.352–381.

Biographical notes: Ambjörn Naeve has a background in Mathematics and Computer Science and received his PhD in Computer Science from KTH in 1993. He is presently coordinator of research on interactive learning environments and semantic web at CID/KTH, and Head of the *Knowledge Management Research* group. He is also a well-known industry consultant with extensive experience in UML modelling for software engineering and business applications. Ambjörn Naeve has invented the concept browser and has developed a modelling technique called *Unified Language Modelling*, which is specially designed to depict conceptual relationships in a linguistically coherent way.

Mikael Nilsson has a background in Mathematics and Computer Science. He is currently a fourth year PhD student in the Swedish National Research School of Mathematics Education with a focus on ICT-enhanced mathematics education based on semantic web technology. He is a major designer and developer of the infrastructure *Edutella*, the frameworks *SCAM* and *SHAME*, and the knowledge management tool *Conzilla*. He is heavily involved in several international standardisation efforts within technology-enhanced learning, and he is presently coordinating the development of the RDF binding for the IEEE-LOM metadata standard.

Matthias Palmér has a background in Mathematics and Computer Science. He is currently a fourth year PhD student in Computer Science at CID/KTH with a focus on technology enhanced learning environments and the semantic web. He is a major designer and developer of the infrastructure *Edutella*, the frameworks *SCAM* and *SHAME*, and the knowledge management tool *Conzilla*. In his research, these projects form the basis for a prototype implementation of a learner-centric educational architecture based on open source and emerging international ICT standards. He is therefore involved in standardisation efforts around *LOM* and *IMS Content Packaging*.

Fredrik Paulsson has a background in Computer Science. Since 1999 he is Technical Advisor at the National Agency for School Improvement and since 2001 he is an industrial PhD student at CID/KTH. Since 2002 he has been coordinator for e-learning at the Centre for IT in Northern Sweden. Between 1993 and1996 he taught at LITU/Umeå University, and between 1996 and 1999 he was responsible for projects and development there. He is the technical coordinator of the *SCAM* project, and he is a member of several steering – and reference groups for technology enhanced learning at both the national and the European level.

1 Introduction

Today educational technologies are reaching a state that allows interoperability and reuse of learning resources. The underlying techniques rely heavily on the standards movement for metadata representation. On top of this, a number of monolithic reference platforms are being developed with the aim to ease application development. However, we do not think this approach is flexible enough to embrace future learning techniques (Palmér et al., 2001). In contrast, we suggest a learning environment based on open source and open international ICT standards, where educational services can be developed and exchanged between as well as within systems. We call this learning environment a *Public e-Learning Platform* (PeLP).

In this paper we present a number of contributions to the PeLP in the form of *infrastructure*, *architecture*, *frameworks* and *tools*. They are based on the technology for the emerging next generation internet – the so-called *semantic web* (http://www.semanticweb.org).

Specifically we present:

- *the Edutella infrastructure*: a democratic (peer-to-peer) network infrastructure for search and retrieval of information about learning resources
- *the knowledge manifold architecture*: *a*n information architecture that consists of a number of linked conceptual information landscapes (context-maps), whose concepts can be filled with content
- *the Conzilla concept browser*: a knowledge management tool that supports the construction, navigation, annotation and presentation of the information in a knowledge manifold

- *the SCAM framework*: a framework that helps applications to store and share information about learning resources
- *the SHAME framework*: an editor framework that supports an evolving annotation process of learning resources in a way that enables the growth of an 'ecosystem' of quality metadata
- *the Formulator (or SHAMEditorEditor)*: a tool for editing metadata editors that is built on top of the SHAME framework
- *the SCAM portfolio (or ConFolio)*: an e-portfolio system that is built on top of SCAM, SHAME and Edutella, and which supports collaborative and reflective learning techniques
- *the VWE composer*: an environment for composing learning resources and building customised learning modules

2 Background

2.1 The unified language modelling technique

Unified Language Modelling is a context-mapping technique, which has been developed by Naeve (1997, 1999) during the past decade. It is designed to visually represent a verbal description of a subject domain in a coherent way. Today, the ULM technique is based on the *Unified Modelling Language* (Rumbaugh et al., 1999), (http://www.uml.org), which is a de facto industry standard for systems modelling.

Figure 1 The basic verbal/visual correspondence of Unified Language Modelling

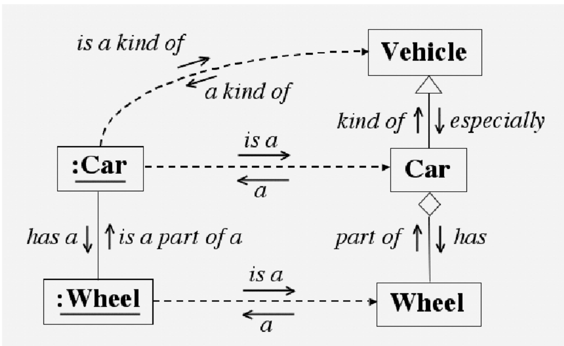

In ULM the resulting context-maps have a clearly defined and verbally coherent visual semantics, which makes it easy to cognitively integrate the conceptual relations and achieve a clear overview of the context. Moreover, making the context visually explicit provides important support for the conceptual calibration activities that form an integral part of the learning process. The ULM verbal-to-visual contextual representation technique has a crucial advantage in comparison with similar techniques such as *concept maps* (http://www.graphic.org/concept.html) or *topicmaps* (http://www.topicmaps.org), which have to rely on purely verbal semantics in order to convey their conceptual relationships.

2.2 The modelling theory of David Hestenes

We believe that the activity of building and navigating the 'conceptual landscapes' expressed by context-maps is of fundamental importance for the quality of the learning process. However, we have only recently begun to gather empirical evidence for this hypothesis, so in order to motivate it, we will briefly describe the basic ideas that underlie the modelling theory of David Hestenes – a well-known physicist and physics education researcher. Although he has applied his modelling theory mainly to physics education (http://modelingnts.la.asu.edu/html/Modeling.html) – and over the past two decades achieved striking results within this field – we share his belief that his theory is applicable to learning in general.

Hestenes uses the term 'conceptual learning' for the type of learning that is the opposite of 'rote learning'. The following is a brief presentation[1] of Hestenes' five general principles of conceptual learning that he has incorporated into his instructional theory and applied repeatedly in the design of instruction.

- *Conceptual learning is a creative act*. This is the crux of the so-called constructivist revolution in education, most succinctly captured in Piaget's maxim: 'To understand is to invent!' Its meaning is best conveyed by an example: for a student to learn Newtonian physics is a creative act comparable to Newton's original invention. The main difference is that the student has stronger hints than Newton did.[2]
- *Conceptual learning is systemic*. This means that concepts derive their meaning from their place in a coherent conceptual system. For example, the Newtonian concept of force is a multidimensional concept that derives its meaning from the whole Newtonian system. Consequently, instruction that promotes coordinated use of Newton's laws should be more effective than a piecemeal approach that concentrates on teaching each of Newton's laws separately.
- *Conceptual learning depends on context*. This includes social and intellectual context. It follows that a central problem in the design of instruction is to create a learning environment that optimises the learner's opportunities for systemic learning of targeted concepts. The context for scientific research is equally important, and it is relevant to the organisation and management of research teams and institutes.
- *The quality of learning depends on the conceptual tools*. The quality of learning is critically dependent on conceptual tools at the learner's command. The design of tools to optimise learning is therefore an important subject for educational research (Hestenes, 1995).
- *Expert learning requires critical feedback*. Expert learning requires deliberate practice with critical feedback. There is substantial evidence that practice does not significantly improve intellectual performance unless it is guided by critical feedback and *deliberate* attempts to improve. Students waste an enormous amount of time in rote study that does not satisfy this principle.

The textual interpretation of Figure 2 is the following: Conceptual Learning *is a* Creative Act. Conceptual Learning *is* (*a*) Systemic. Conceptual Learning *depends on* Social Context and Intellectual Context. Conceptual Learning *depends on* Learning Tools, *especially* Modelling Tools. Conceptual Learning *depends on* Practice, which *consists of* (= *has*) Deliberate Improvement Attempts and Critical Feedback.

Figure 2 A ULM context-map of Hestenes' five learning principles

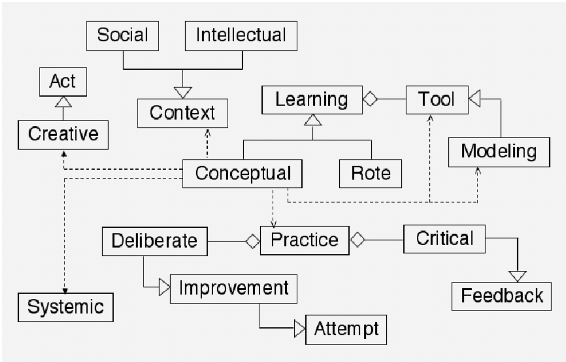

In his response (Hestenes, 2002) to the Oersted medal reward in 2002, Hestenes writes:

> "I believe that all five principles are essential to effective learning and instructional design, though they are seldom invoked explicitly, and many efforts at educational reform founder because of insufficient attention to one or more of them."
>
> "I see the five Learning Principles as equally applicable to the conduct of research and to the design of instruction. They support the popular goal of teaching the student to think like a scientist."

2.3 The knowledge manifold information architecture

Naeve (1997, 1999, 2001b) has invented an information architecture called a *knowledge manifold*, which highlights the complementarity of *context* and *content* and supports a variety of different strategies for context-dependent presentation and suppression of information. It consists of a number of linked information landscapes (contexts), where one can navigate, search for, annotate and present all kinds of electronically stored information.

The KMR group (http://kmr.nada.kth.se) at CID (http://cid.nada.kth.se/en) is making use of the knowledge manifold architecture in the construction of interactive learning environments that enable a learner-centric, interest-oriented form of 'knowledge-pull', and which support inquiry-based and personalisable forms of networked learning. An important design goal for these learning environments is to support the transformation of the teaching role – away from the traditional 'knowledge filter' towards more of a 'knowledge coach', i.e., away from 'teaching you what I know' and towards 'helping you to find out more about the things that you are interested in'. We believe that this provides a way to create sustainable solutions to the current educational crisis (Douglas, 1991; League of World Universities, 1993; Learn in Freedom).

2.3.1 Fundamental pedagogical assumptions

The knowledge/learning manifold educational architecture is based on the following fundamental pedagogical principles:

- Nobody can teach you anything. A good teacher can inspire you to learn.
- Your motivation to learn is based on the experience of subject excitement and faith in your learning capacity from live teachers.
- Your learning quality is enhanced by taking control of your own learning process.
- No 'problematic' questions can be answered in an automated way. In fact, it is precisely when your questions break the pre-programmed structure that the deeper part of your learning process begins.

2.3.2 Basic structure

A knowledge manifold consists of a number of linked *knowledge patches* – each maintained by a custodian called a *knowledge gardener*. A knowledge patch in turn consists of a set of resource components[3] that are tied together with *context-maps* that represent the corresponding conceptual model of the subject domain. Such context-maps are preferably constructed using the ULM technique described above.

A knowledge manifold has the following major characteristics:

- it can be regarded as a *knowledge patchwork*, with a number of *linked knowledge patches*, each maintained by its own *knowledge gardener*
- it gives the learners the opportunity to ask questions and search for *live, certified knowledge sources* to discuss them with
- it has access to distributed archives of *resource components*, which are described by an 'ecosystem'[4] of evolving metadata annotations
- it allows 'knowledge composers' to construct customised learning modules by *composing resource components*
- it makes use of *conceptual modelling* in order to construct context-maps whose concepts (and concept-relations) can be filled with content[5]
- it contains a *concept browser* that lets the user navigate the context-maps and view their content filtered by a dynamically configurable set of context-dependent aspects.

2.3.3 The seven different knowledge roles

When used for learning purposes, the KM architecture supports the following seven different knowledge roles:

- the *knowledge cartographer*, who constructs and maintains context-maps
- the *knowledge librarian*, who fills context maps with content-components
- the *knowledge composer*, who constructs customised learning modules
- the *knowledge coach*, who cultivates questions
- the *knowledge preacher*, who provides live answers
- the *knowledge plumber*, who directs questions to appropriate preachers
- the *knowledge mentor*, who is a role model and supports self-reflection.

It is fundamentally important that all these roles should be available to both teachers and learners.[6] 'You learn as long as you are teaching' is the pedagogical principle at work here.

2.4 The conceptual browsing of a knowledge manifold

A *concept browser* (Naeve, 1999, 2001a) is a form of knowledge management tool that is tailored to the knowledge manifold architecture and which enables the effective organisation, annotation, navigation and presentation of its information in a number of different ways. Most notably, by supporting the separation between conceptual *context* and *content*, it is possible to study the content without losing overview of the context. A concept browser facilitates navigation between all the different contexts where a given concept occurs. Moreover, each concept can be equipped with a set of resources (content-components) that can be presented in various ways by filtering them through a set of context-dependent aspects.

A concept browser conforms to the following major design principles:

- Separate the *content* of a concept from its *contexts*. This supports the reuse of content across different contexts.
- Describe each context in terms of a *context-map*, preferably by using the Unified Language Modelling technique described above.
- Support contextual navigation on each concept by enabling the direct switch from its presently displayed context into any one of the other contexts where this concept appears.
- Assign appropriate resources as the content-components of each concept.
- Embed the resources in a 'metadata ecosystem', by supporting multiple (subjective) descriptions and evolving annotations of each resource expressed in a standardised (semantic) metadata language.
- Allow metadata-based filtering of the resources through context-dependent aspect filters. This enables the presentation of the content in a way that depends on the context.
- Enable the unpacking and browsing of context-maps that have been stored as content-components.

3 The emerging public e-learning platform

3.1 Overview

In this chapter we will describe the structure of this Public e-Learning Platform (PeLP) and illustrate how you can work with its infrastructure, architecture and tools – both as an application developer and as an end-user. We will discuss the emerging next generation internet – the (machine) *semantic*[7] *web* – and introduce a more human-understandable interface to it, which we call the *conceptual web*. The knowledge manifold architecture is inherent in the various knowledge roles that are associated with the respective tools below.

Specifically, we will discuss the infrastructure, architecture, frameworks and tools of the PeLP and show how they can:

Help application developers on the semantic web to design and implement:

- providers and consumers of metadata through the Edutella infrastructure
- storage systems for metadata built on top of the SCAM framework
- flexible editors for metadata built on top of the SHAME framework.

Help end-users on the semantic web to:[8]

- build metadata archives in a SCAM portfolio (librarian)
- expose metadata archives on Edutella through a SCAM provider (librarian)
- edit metadata in a SHAMEditor (cartographer, librarian)
- edit a SHAMEditor in a SHAMEditorEditor (cartographer, librarian)
- query and search for metadata on Edutella through a SHAME consumer (librarian)
- combine resources and build learning modules in a VWE composer (composer).

Help application developers on the conceptual web to:

- create context-maps that visualise conceptual *models*[9] through the ULM technique
- connect the context-maps through the knowledge manifold architecture.

Help end-users on the conceptual web to:

- construct, browse and edit knowledge manifolds in the Conzilla concept browser (cartographer, librarian, composer).

Figure 3 Overview of the KMR contributions to a PeLP

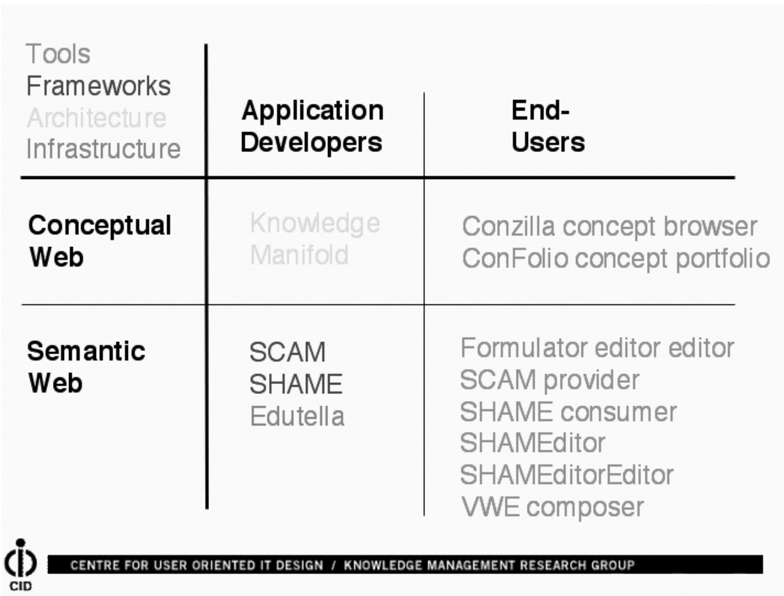

The KMR group is presently coordinating a collaborative effort that involves the Swedish National Agency for School Improvement (MSU), the Swedish Educational Broadcasting Company (UR), and the Swedish National Centre for Flexible Learning (CFL). These three major Swedish public service e-learning players have teamed up and are now jointly contributing to the PeLP. The Soft infrastructure for IT in education project of MSU (http://mjukis.skolutveckling.se), the Digital Media Library of UR (http://www.ur.se/mb), and the Learning Resource Centre of CFL (http://larresurs.cfl.se) are three of the important stakeholders projects in the PeLP.

3.2 The semantic web

The W3C[10] has created an initiative called the *semantic web* (http://www.semanticweb.org), which embodies the vision of the next generation of the internet. The stated goal of the semantic web is to enable machine understanding of information about web resources, i.e., metadata. The rationale behind the development of the semantic web has been that deriving meaning from contemporary HTML or other web resources is nearly impossible due to the lack of a common metadata framework for describing resources.

The technical basis for the semantic web is a metadata language called RDF (http://www.w3.org/RDF), which makes it possible for anyone to 'state anything about anything' in a way that is machine-understandable.[11] In fact, most resource descriptions today are in the form of natural language text embedded in HTML. While such semantic descriptions are meaningful only to the human reader, the semantic web will provide such descriptions in machine-readable format.

Through the emergence of the semantic web, the metadata has acquired the potential to become just as distributed as the data it describes – while still remaining just as 'searchable and combinable' as if it resided inside a single database. This metadata decentralisation process – which is enabled by the representational power of RDF – is bound to have a profound impact on the design and use of ICT-supported learning environments in general (Nilsson, 2001).

Metadata will no longer be restricted to something 'objective' that has to be downloaded from some central server. On the contrary, metadata will be allowed to consist of subjective views of resources that are distributed and shared in contexts that can evolve dynamically (Nilsson et al., 2002). In support of such requirements, our learning framework consists of a combination of semantic web techniques and peer-to-peer services for search, retrieval, publication, replication and mapping of metadata.

3.3 The metadata ecosystem

A popular – but misleading – view of metadata is that it is something you produce once, often when you publish your document or resource, and which remains with the resource for its lifetime.[12] This is the way metadata is implemented in most systems that support it. This conception is related to the idea of metadata as being authoritative, objective information consisting of facts that do not change. The problem with implementing metadata support in this way is that it efficiently hinders subjective opinions and context-dependent metadata.

Instead, metadata needs to be handled as a continuous work in progress, where updating and modifying descriptions is a natural part of the metadata publishing process. Treating metadata as a continuous work in progress and allowing subjective metadata leads to a new view of metadata. Metadata is information that evolves, constantly subject to updates and modifications. Competition between descriptions is encouraged, and thanks to RDF, different kinds and layers of context-specific metadata can always be added by others when the need arises. Any piece of RDF metadata forms part of a global network of information, where anyone has the capability of adding metadata to any resource.

Metadata for a resource need not be contained in a single RDF document. Translations might be administrated separately, different categories of metadata might be separated, and additional information may be provided by other sources. Consensus building becomes a natural part of metadata management, and metadata can form part of the ongoing scientific discourse. The result is the global *metadata ecosystem* described in (Nilsson et al., 2002), a place where metadata can flourish and cross-fertilise, where it can evolve and be reused in new and unanticipated contexts, and where everyone is allowed to participate. This provides support for the conceptual calibration process in a bottom up fashion, which builds consensus in the same way as it is achieved between people.

It is important to realise that metadata is not only for machine consumption. In the end, computers are a medium for human-to-human communication, and conceptual metadata that is understandable for both the human and the machine is a strategically important part of this communication process.

3.4 The Edutella infrastructure

Within the *Wallenberg Global Learning Network* (www.wgln.org), the KMR group is participating in an international collaboration project called *PADLR*[13] (http://www.learninglab.de/padlr/index.html) that involves a number of different institutes and research groups – notably *L3S Research Center*[14] (http:// www.learninglab.de), *Uppsala Learning Lab* (www.ull.uu.se), *Uppsala DataBase Laboratory* (http://user.it.uu.se/~udbl), and *AIFB* (http://www.aifb.uni-karlsruhe.de/ english) at the University of Karlsruhe. The driving vision of the PADLR project is a learning web infrastructure, which will make it possible to exchange/annotate/organise and personalise/navigate/use/reuse modular learning resources, supporting a variety of courses, disciplines and educational stakeholders. The PADLR project has initiated the development of *Edutella* (Decker et al., 2001; Wilson, 2001; Nejdl et al., 2002), an infrastructure and a search service for a peer-to-peer network that facilitates the exchange of metadata about educational resources on the semantic web. The Edutella project (http:// edutella.jxta.org) also includes a number of other participants, such as *UNIVERSAL* (http://www.ist-universal.org) and *EducaNext* (http://www.educanext.org), and it is still expanding. Edutella will also play an important role within the recently established *Prolearn*[15] *network* (http://www.prolearn-project.org), which is a network of excellence within the sixth framework programme of the European Union.

Much of the current work in e-learning technology targets learning objects stored in LMS servers and other centralised structures, often of large proportions. Even though standards such as IEEE-LOM (http://ltsc.ieee.org/wg12) increase the interoperability of such systems, they are still mostly isolated information islands.[16] Cross-search of repositories is not a reality. In fact, it has been said that the web is still in the 'hunter-gatherer phase' with respect to searching. This is certainly true for learning objects. We have not yet reached the goal of a global e-learning community.

In addition, many institutions are reluctant to give up control over their learning resources. This is troubling many central-server-based approaches to learning resource sharing, often designed as e-learning portals. For this and other reasons, such portals are costly and difficult to maintain.

In contrast, Edutella takes a different route (Nilsson et al., 2004). It is a piece in an e-learning infrastructure with a decentralised vision. By encouraging small-scale content repositories, anyone can participate in the exchange and annotation of e-learning resources. By allowing anyone to participate, the learner is given more control over the learning process, leading us towards the vision of a learner-centric educational architecture.

The envisioned Edutella services will include *searching*, *mapping* and *replication*. Searches will be routed to anyone who has registered a matching answering capability. Mapping will enable translation between different metadata representations (schemas), something that will allow very flexible reuse of information. An application will not need to adapt to competing or more capable schemas, because these schemas can be mapped to something that the application already understands. Replication will allow metadata about learning resources to be spread across the network, which will simplify the discovery of the corresponding resources.

The search service of Edutella is currently functional, while the other services are still being researched. At the time of writing of this paper, Edutella is being tested in a number of national and international projects, mainly based in Europe.

3.4.1 Edutella technology

Edutella connects a multitude of highly heterogeneous peers,[17] and the goal of the Edutella project is to make the distributed nature of the Edutella services (e.g., repository search) completely transparent to Edutella clients.

The first building block of Edutella is an open-source peer-to-peer technology called JXTA (http://www.jxta.org) initiated by Sun Microsystems. JXTA is a generic peer-to-peer protocol, which is designed for use in many diverse kinds of peer-to-peer applications, and which is focused on interoperability, platform independence and ubiquity.

The second building block of Edutella is RDF, which is a highly sophisticated framework for expressing metadata of any kind, containing facilities for combining resource descriptions using different vocabularies from different sources.

To show the kind of queries that Edutella can manage, consider the following Edutella query, constructed in the Conzilla concept browser (www.conzilla.org) described in Section 3.10.

In Figure 4, *X* represents the resource we are looking for. The arcs are properties of that resource. In plain English, the query asks for (in counter-clockwise order):

> "All Scientific Works on the subject of Politics, having Lebanon as subject or keyword, with a title (*Y*), written in English, German or French, created or contributed to by a Person (*Z*), employed at a University, and created after 1980."[18]

Figure 4 An RDF-based Edutella query (top) and its conzilla interface (bottom)

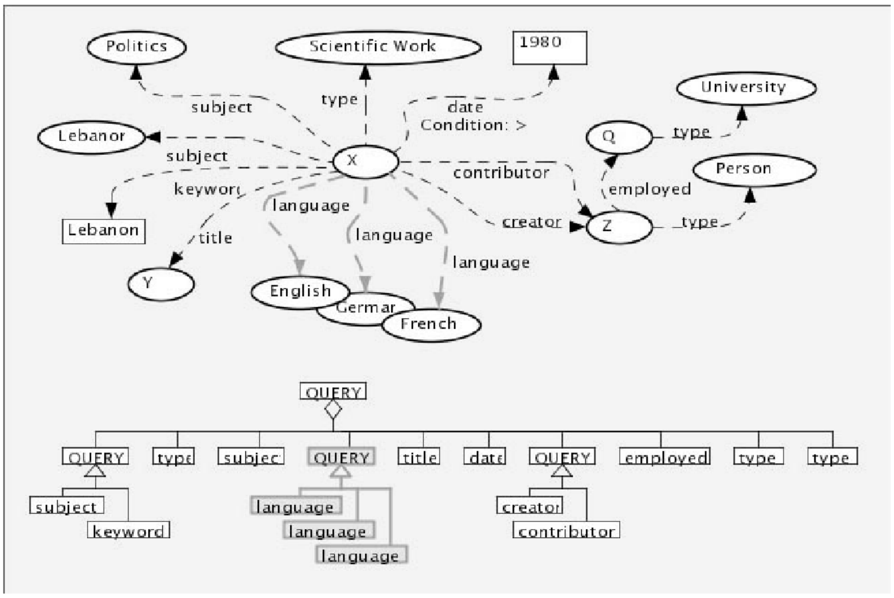

Edutella takes queries of the above complexity, distributes them to peers that have declared themselves capable of answering this type of query, collects the answers and returns them to the originator. It is possible that parts of the answers are located on different peers. In the example, the university employee information is perhaps not located on the same server as the resource metadata, but Edutella is able to handle these kinds of situations in a transparent manner.

3.4.2 Nodes in an Edutella network

When looking for information on Edutella, your question will be routed to peers that can answer your query, and they will return matching results to you. In order to be a provider, all that is required is that you are able to answer questions formulated in the Edutella query language – so any kind of information source can be given an Edutella interface.

Examples of consumers that could use Edutella to find information could be:

- a search tool in an LMS system that uses Edutella to get answers
- a generic, self-contained search tool, such as Conzilla or SHAME consumer, or a domain-specific search tool such as the SWEBOK consumer (http://edutella.jxta.org/downloads)
- an end-user application that uses Edutella to enhance the user experience with metadata information (such as e.g., 'related material')
- an augmented-reality system that displays and uses metadata for objects in three-dimensional space (real or virtual)
- a web portal that includes an Edutella search interface

- a mobile device (PDA, cell phone, etc.) that gathers information from Edutella to enhance your stay in Rome
- a smart software agent that gathers relevant information from Edutella to help you construct your personal learning environment
- a crawler or push-based system such as CourseWare Watchdog (http://www.aifb.uni-karlsruhe.de/Publikationen/showPublikationen?id_db=51), which uses Edutella as an additional information source.

3.5 The SCAM framework

The *Standardised Contextualised Access to Metadata* (SCAM) *framework* (Paulsson and Naeve, 2003; Palmér et al., 2004), (http://scam.sourceforge.net) constitutes a general basis for constructing standardised archives for digital information about learning resources. This means that the use of international learning technology standards[19] (as well as other technology standards) is most essential. The work is based on the assumption that the exclusive adherence to international metadata standards is prevented by the diversity of the applications that will be built on top of the SCAM framework. At the same time, a great part of the implementation is similar for most applications in this domain. Hence, a common basis greatly increases the effectiveness by enhancing reuse as well as by hiding the complex implementation details – thereby providing a higher abstraction level for the average application developer.

3.5.1 Metadata and organisation in the SCAM framework

One of the most important missions of the SCAM framework is to serve as a *metadata repository* for learning resources. The resources themselves may be distributed and referred to by URIs.[20] The SCAM framework is not limited to a specific metadata set – such as LOM – but instead it relies on RDF as a general metadata format. Since there is a newly released RDF binding for LOM/IMS Meta-data (Nilsson et al., 2003),[21] we do not lose any expressiveness regarding learning resources.

The SCAM framework also supports different – and sometimes even conflicting – annotations about a resource by providing separate contexts for metadata within a single SCAM repository. Each context represents a single metadata source – typically an individual or an organisation. By making use of contexts, metadata is made more accessible and manageable to humans. In this way SCAM is aiding the development of the conceptual web, which is discussed in Section 3.9.

In order to support the re-combination of resource components, the SCAM framework provides an RDF-based version of IMS Content Packaging (http://imsglobal.org/content/packaging).[22] This standard allows a learning resource component to be used in several different contexts, each of which is described with its own context-specific metadata. An important design goal of SCAM has been to allow interoperability with other IMS-compliant systems (http://www.imsproject.org).

SCAM also implements an Edutella peer, the *SCAM provider*, which allows a suitable portion of the stored metadata to be searchable over the Edutella network.

3.6 The SHAME framework

The SCAM framework manages many different kinds of metadata as long as they are expressed in RDF. To make full use of the flexibility of the SCAM framework, an equally capable front-end for presenting and editing metadata is needed (Kraan, 2003).

The *Standardised Hyper-Adaptable Metadata Editor* (SHAME) (http://kmr.nada.kth.se/shame) is a *framework* for RDF-based metadata editors rather than a single editor (http://kmr.nada.kth.se/imsevimse).[23] On top of the SHAME framework one can construct specific *SHAMEditors*, which are configured for specific purposes and specific metadata sets. A SHAMEditor defines not only which metadata to edit, but also which pre-defined values (taxonomies) that should show up in the drop-down menus,[24] which order of presentation that should be used, which metadata that should be suppressed, which metadata that is allowed to occur several times, and which metadata that should be restricted to specific data types.

Another feature of the SHAME framework is that the design of a SHAMEditor is not directed toward a specific GUI environment such as a specific programming environment or a specific web interface. In contrast, a SHAMEditor describes an abstract editor, which is used by a suitable GUI factory within the SHAME framework in order to generate a concrete editor. We have produced two such proof-of-concept implementations of concrete editors, one using JSP for a web interface and the other using pure Java (Swing).

For practical reasons, a SHAMEditor is expressed in RDF and is therefore rather hard to edit by hand. Hence there is a need for a SHAMEditor editor. Obviously it is nice to use SHAME itself for that purpose. We call this editor *Formulator* or SHAME^2 (=SHAME-squared), which expands to *SHAMEditorEditor*.

Another offered functionality is to use a SHAMEditor in order to generate Edutella queries. By using a partially completed form as a matching pattern, this form can be used as a query on the Edutella network. Since SHAME uses a formalism that is very close to the Edutella query language, there is very little translation involved. This can be viewed as using the editor 'in reverse' and therefore we call this query method SHAME^{-1} (=SHAME-inverse). It allows experts to write queries that users can choose among and specialise further by filling out a simple form. The result obtained from executing a query can be displayed via a presentation defined by the same SHAMEditor that specified the query. This functionality is provided by the *SHAME consumer*, which is a stand-alone tool that provides a set of general queries that users can fill in and execute on the Edutella network.

3.7 The SCAM portfolio/ConFolio

One of the many possible user-interfaces to SCAM is a digital portfolio (e-portfolio) – a personal online repository of information, which is used in e-learning activities by both teachers and learners for storage and publishing of information. We have designed the SCAM portfolio so that it uses RDF descriptions of both metadata and structure, making use of the IMS metadata and content packaging standards for that purpose. As discussed above, some of the fundamental problems and common misconceptions regarding metadata is that it is objective, static and has logically defined semantics. Since this is not really true, especially not for learning resources, we need a mechanism in the portfolios for supporting a metadata ecosystem of dynamically evolving metadata over multiple vocabularies and taxonomies. Since SCAM is agnostic to what kind of metadata it stores,[25] the challenge is to provide flexible interfaces to the metadata, which is where the SHAME framework comes in. In the SCAM portfolio, we use SHAMEditors to specify the interface for both editing, presenting and searching for metadata.

Together with the facilities for metadata storage, the SCAM portfolio also provides annotated storage for actual learning resources, or annotated links to such resources if they already exist on the web. Hence, when equipped with an Edutella peer interface, a SCAM portfolio, also called *ConFolio*, becomes a content management system allowing not only the publishing of documents, but also the dissemination of metadata about documents and the structure of courses, as well as subjective annotations of online resources. Taken together, these properties effectively support collaborative and reflective learning techniques. SCAM portfolios have been introduced to teachers and students at KTH and Uppsala with promising results (Blomqvist et al., 2003).

3.8 The VWE composer

The *Virtual Workspace Environment* (*VWE*) *composer* (www.vwe.nu) is a component-based virtual learning environment framework, which is designed to support the construction of customisable learning environments by the composition of components (VWE tools). VWE tools are actually software components that provide the functionality for a VWE workspace.[26]

The basic idea behind VWE is to make it possible for teachers and learners to create and administrate their own learning environments in the shape of *workspaces* – based on the requirements for a specific activity and learning situation. A workspace is created by choosing and combining the desired functionality from the available VWE tools. This process works in a way that is similar to a LEGO building kit. The tools in a workspace may provide any kind of functionality, from a simple chat application to more advanced applications such as word-processing, spreadsheet, simulation or videoconferencing, all depending on the specific requirements for a certain activity and the available tools. The purpose is to enable teachers and students to control their own tools and provide them with a technology that is adaptable to different learning situations.

At the heart of VWE there are five *services* that handle common tasks as well as communication within the system. These services are *the user service, the tool service, the message service, the file service,* and *the compatibility service*. When a VWE session is started, a small component called the *kernel* is downloaded to the client (currently a web browser). The kernel handles the interaction between the web environment and the VWE services. Client-tools are downloaded at run-time, when they are needed. In fact, VWE can be compared to a small configurable operating system for the web, where the tools are equivalent to the installed applications on a desktop computer. This allows the user to get access to his or her personal learning environment and *WebTop* from anywhere. VWE is most suitable for use with broadband connections, but may also be used with lower bandwidth. The performance depends critically on the chosen tools.

VWE has a SCAM/RDF backend that enables semantic tool interoperation. However, the main reason for the implementation of a SCAM/RDF backend is to enable SCAM to function both as a distributed toolbox for VWE tools and as a distributed file system for VWE users. This makes it possible to organise VWE tools and user files using semantic metadata with all its inherent possibilities.

VWE supports features like multiple languages, change of the conceptual model of the system, different look-and-feel, etc, all in order to enhance the flexibility and the adaptability of the environment. VWE is mainly implemented using Enterprise Java technology, which makes it a scalable platform for building interactive learning environments. VWE makes extensive use of open standards – both general technology standards and learning technology standards – in order to enhance the integration and interaction with other systems.

VWE runs under the Java-2 environment, which makes it platform independent. There is no need for any additional installations at the client. This makes it possible to run VWE in an ordinary web browser, or any other client that supports Java, such as a thin client, a cellular phone, a PDA, etc. Most networked Java application can be adapted and integrated with the VWE via a simple API. VWE supports the use of web services (http://www.w3.org/2002/ws) for learning by its use of standards such as SOAP (http://www.w3.org/TR/SOAP). It is also possible to use non-java tools through the use of the web service framework.

3.9 The conceptual web as a knowledge manifold

Expressing metadata that is machine-understandable is a step in the right direction. However, in order to harness the powers of the semantic web, it needs a 'conceptual interface' that is more comprehensible for humans. We see this conceptual interface as a knowledge manifold,[27] which we call the *conceptual web* (Naeve et al., 2001) or the *human-semantic web*. It serves as a human-understandable 'front-end' that connects to the machine-understandable 'back-end' of the (machine) semantic web.

An important feature of the human-semantic web is the ability to collect metadata from various sources into suitable contexts. These contexts can then be presented in various graphical user interfaces such as adorned forms, through diagrammatic languages such as UML, etc. Such contexts are also usable for the human management of evolving metadata.

The most important contexts of the conceptual web are described by context-maps that are constructed by *conceptual modelling* and connected through the knowledge manifold architecture. This provides a 'conceptual information atlas' of connected context-maps with human-understandable semantics for both abstract ideas and concrete resources. As discussed in Section 2.1, for the conceptual modelling we make use of the Unified Language Modelling technique (based on UML) that is tailored to support the visualisation of how we speak about things. UML provides a well-proven and standardised modelling vocabulary with clearly defined visual semantics of the relationships between the occurring concepts.

Combining the human semantics of UML with the machine semantics of RDF enables more efficient and user-friendly forms of human-computer interaction. Within the e-learning field, the conceptual web will support the mixture of human- and machine semantics that is needed for efficient construction and use of modular and personalised learning environments based on retrieval and reuse of relevant learning resources.

3.10 The Conzilla conceptual web browser

As described in Section 2.4, the basic navigation and presentation tool of a knowledge manifold – i.e., of the conceptual web – is the concept browser. This tool allows the user to browse conceptual contexts in the form of context-maps with rich annotations.

Our incarnation of a concept browser is called *Conzilla* (www.conzilla.org) and has been developed by the KMR group as an open source project over the last six years.

3.10.1 Conzilla 1: the first prototype

The first version, Conzilla 1, is based on XML (Nilsson, 2000), supports multiple languages and conforms to the IEEE-LOM metadata standard (http://ltsc.ieee.org/wg12). Conzilla 1 has proved to be a valuable tool for providing an overview of complex web-based material. It gives the user a clear overview of the subject area (=context), while at the same time allowing the exploration of its various forms of content. Incorporating web resources as content is done by associating concepts with occurrences in resources. This has the important benefit of a clear and browsable visual overview of the context while viewing the content in, for example, an ordinary web-browser.

Using Conzilla 1, several different knowledge manifolds have been constructed by the KMR group, e.g., within the fields of mathematics (Naeve and Nilsson, 2004), e-administration (Naeve and Selg, 2001), ICT-standardisation, and interoperability between different systems for e-commerce (http://www.ecimf.org).

A glimpse of the virtual mathematics laboratorium

In order to convey an idea of what a Conzilla-based knowledge manifold can look like, we will now take a brief look at the *Virtual Mathematics Laboratorium* (VML), which is a mathematical knowledge manifold. The mathematical concepts are described with metadata and filled with content-components according to the general design principles for knowledge manifolds that have been outlined above. Moreover, *aspect filters* (Pettersson, 2000) allow the selective viewing of the content-components – based on different aspects and levels of difficulty.[28] Currently, the VML contains more than 500 different interactive mathematics components (learning objects).[29]

The idea is that learners should be able to browse through the conceptual information landscapes of the VML in order to get an overview of the corresponding mathematical concepts and their relations – if necessary by inspecting their respective metadata annotations. In addition, both teachers and learners should be able to find resource-components that cover the relevant aspects of the topics they are interested in at the appropriate level of complexity. These components can then be combined into customised learning modules using the VWE composer.

Figure 5 shows a context-map of the most common kinds of numbers in mathematics. Moreover, the map also shows the relationships between these kinds of numbers. More precisely, the map show [30] that a *Natural Number* is a kind of *Integer Number,* which is a kind of *Rational Number,* which is a kind of *Real Number,* which is a kind of *Complex Number*. In Figure 5, the metadata on these relationships have been exposed, which gives more detailed information about how these different kinds of numbers are related. It is important to observe that this metadata belongs to the *relations* (=arrows) between the boxes that represent the different number-types. The metadata has been exposed by pointing to the respective relations (one-by-one) and right clicking each one and selecting the *Info* command from the pop-up menu that is visible in Figures 6 and 7. In this mode, Conzilla works as a nonlinear presentation tool – as opposed to e.g., PowerPoint, where the presentations are linear – i.e., totally ordered with no possible changes at 'run-time'.

Figure 5 Different kinds of numbers with exposed metadata describing their relationships

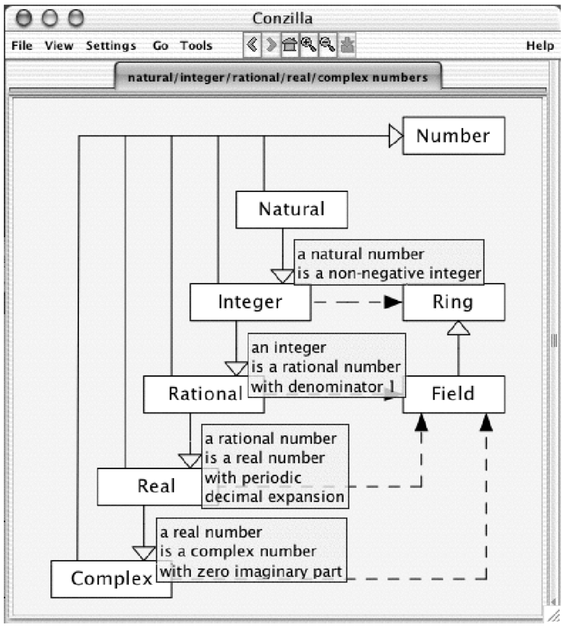

Figure 6 Right clicking on 'Complex' and choosing 'View' opens the content-window

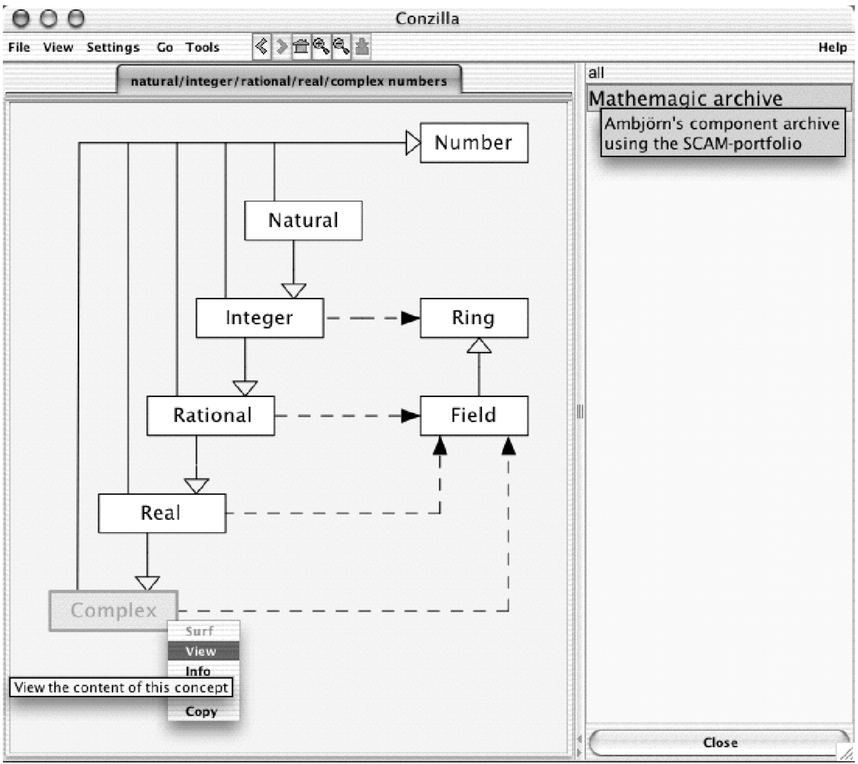

Figure 7 Clicking on 'Mathemagic archive' opens the Complex (number) part of Ambjörn's mathemagic SCAM-portfolio in an ordinary web browser

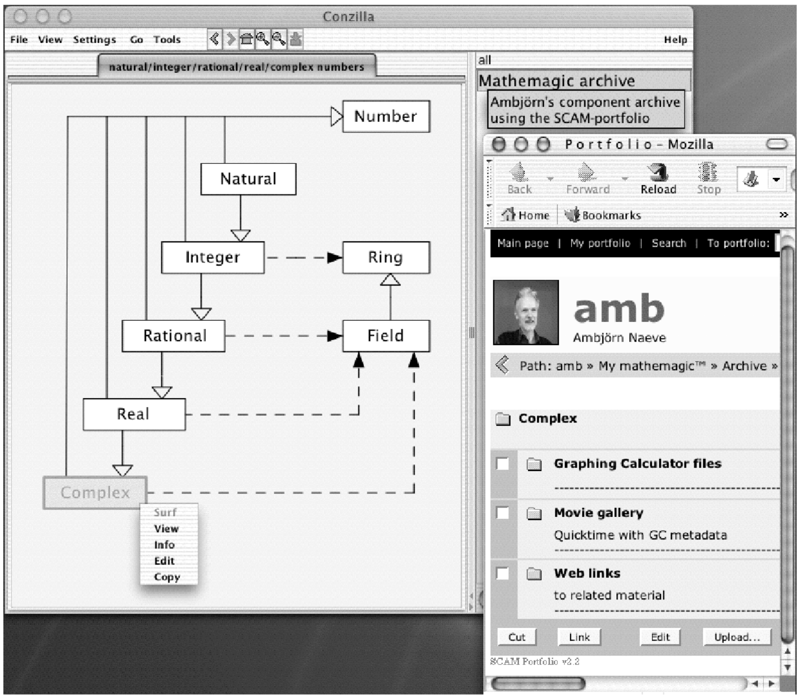

In Figure 6, the concept *Complex* has been right-clicked and the *View* command has been selected from the appearing pop-up menu. This opens the content-window (to the right), which displays a list of content-components that have been associated with the *Complex* concept. In this case there is only one component – named *Mathemagic archive,* and exposing its metadata shows that it is *Ambjörn's component archive using the SCAM-portfolio*. Clicking on *Mathemagic archive* brings up the *Complex* part of this archive in an ordinary web browser – in this case Mozilla – as shown in Figure 7. Now we are ready to dig into the content of this part of the archive, but – in contrast to the experience of ordinary surfing – we can do so without losing the overview of the original context.

Online discussion service

An important type of content-component is an annotated link to a live human knowledge source with an electronically certified identity.[31] This could be a fellow student or a Professor, who has declared a willingness to engage in discussions about the subject where his or her contact information can be found.

This feature supports the KM pedagogical principle[32] that 'no problematic questions can be answered in an automated way'. Hence, an important part of the KM architecture is an *online discussion service*.[33] Instead of discussing teacher-generated questions of a general nature in front of learners that happen to be physically close – but often mentally distant – a teacher should be able to work online and discuss learner-generated questions concerning his or her special interests – questions that come from learners all over the world.[34] This is an important feature of a learner-centric educational architecture.

The closed, layered architecture of the traditional educational systems, especially the lack of contact between universities and schools, is a major problem identified in League of World Universities (1993). The online discussion service makes the KM architecture truly open in terms of educational layers. A learner from elementary-, secondary- or high school could easily get in touch with a lecturer at a university in order to discuss a problem that has eluded a meaningful discussion within the other layers.

3.10.2 Conzilla 2: the conceptual web browser

The next version of our concept browser, Conzilla 2, will be our first full-fledged browser for the conceptual web. It is based on RDF, equipped with a SHAMEditor and has an interface to both Edutella and SCAM, which will allow flexible annotation, cross-network semantic searches, and easy storage of conceptual content.

In Conzilla 2 the full power of visual modelling is combined with the distributivity and universal annotation property of RDF into a hyperlinked web of conceptually clear material – the conceptual web. Together with our form of visually configurable query/search/filter engines, constructed by using ULM and interfacing with Edutella, this will result in a new and pedagogically revolutionary web experience.

Conzilla 2 is close to completion and will be released during 2005. Conzilla[35] is attracting increased attention both on the national and the international level. We are aiming to develop Conzilla into a combined knowledge-, economy- and management-tool by participating in collaborative projects that will expand the capabilities of the programme within the areas of e-learning, e-commerce and e-administration.[36]

4 An e-learning scenario from the near future

Charles has long been curious to learn more about how his TV-set works. Now when he has retired, he decides to give the subject a try. A friend tells him to try a learning community on the internet. Performing a search for communities focusing on this subject area, he is able to locate several available candidates. He chooses to check out one that deals with both radio and television techniques, mainly because it seems to have a history of encouraging newcomers to form learning groups with some guidance from more advanced community members.

He sends this community an e-mail with a short description of his knowledge background and learning interests. A few days later he receives an answer from Sylvia – a community member who is willing to introduce him to the learning practices of this community. Sylvia sets up a meeting with Charles – mediated by a *virtual presence production system* – and they get along so well that Sylvia offers to act as his *knowledge cartographer*, in order to help him to become a *knowledge gardener* and establish his initial *knowledge patch*. She also encourages him to join a newly formed learning group consisting of seven other newcomers with the intention to follow a best practice introductory plan, which Sylvia is familiar with.

Using the *Conzilla* concept browser, Sylvia helps Charles to describe his present understanding of the workings of a TV-set in terms of a number of interconnected *context-maps*. Together these maps form a conceptual model – a vital ingredient in a knowledge patch – that establishes the most important concepts and their relationships in Charles' present view of this field. Starting with the simple question 'what are the needed ingredients in order for a TV-set to work?' and working together, they begin to model a TV-set on a conceptual level, identifying a number of concepts, like the signal – transmitted through a cable or via radio waves – the process that divides the signal into separate channels, the screen, the signal controlling device, the power supply, etc. In this process Charles learns to express his thoughts and ideas in the *Unified Modelling Language* – using the *Unified Language Modelling* technique. On the one hand this helps him to establish a conceptual overview for himself and formulate the questions that he would be prepared to spend energy pursuing. On the other hand, it makes it possible for him to transcend the traditional language barriers and communicate internationally with other learners within the same subject area. As Charles connects the model to other related conceptual models using *semantic web* technologies, his model becomes part of the *conceptual web.* In the next step the group compares their models, using Conzilla to describe similarities and differences. After much discussion, partial consensus is reached and collective work on a common model is initiated. This collective model forms the basics of a *shared knowledge patch*, where their respective contributions can be tracked back to their original knowledge patches.

Gardening his own knowledge patch of context-maps – and collaboratively the patch of the group – Charles learns to supply the concepts with *metadata*, expressing their properties as well as his opinions as to what needs to be further explored about them. He also learns how to provide the concepts of the context-maps with *content-components*, i.e., links to associated relevant resources that he is able to locate on the internet.

After a while, the group members feel that their respective knowledge patches reasonably well express their present knowledge, as well as their opinions on how they would like to expand it. In the next group meeting they therefore contact Caroline, another community member, who has volunteered to help the group as an experienced *knowledge librarian*.

Through Conzilla, Caroline formulates questions and searches for metadata about learning resources on the Edutella network. Since Edutella treats questions as a form of metadata, Caroline can start by searching for questions that have proven efficient in retrieving *related information*. Finding such questions, she edits them in Conzilla during a group session in order to adapt them to the various knowledge backgrounds and learning interests of each member of the group. Then they launch the modified questions in Edutella, and make use of the search results in order to work on the expansion of their individual knowledge patches as well as on the building of 'conceptual bridges' between them.

The collective (=consensus) knowledge patch grows quickly as different group members find their respective interests and contribute with more specific models, examples, simulation programmes, etc. of areas still uncharted. At a certain point, the members feel that they need some human-to-human interaction around some 'problematic' questions that have been raised in the group. Then they get in touch with Shirley, a *certified knowledge-preacher,* whose knowledge profile matches the needs of the group. Using the virtual presence production system, Shirley gives a couple of online lectures on the problematic issues and engages in discussions with the group members. Between themselves the group members also continuously present and discuss possible extensions to their knowledge patches. In this way they gain knowledge more efficiently than they would do by themselves. As the complexity of their knowledge patches grows, the need for *personalised learning modules* increases.

The group posts a request for help with personalisation of learning modules on one of the community bulletin boards, and after a while they get a response from Peter, another community member, who offers his assistance. Peter is an accomplished *knowledge composer* who knows how to deal with the individual needs of each group member. Using the Conzilla concept browser and the *VWE composer*, Peter directs the construction of customised learning modules for each individual group member. The result is a set of interconnected *activity diagrams* (Rumbaugh et al., 1999) that contain related concepts and proposed activities as well as possible contact links to certified live knowledge sources (=preachers).

Charles and the other group members finish after about two months and then most of them continue in other learning groups. Some time later Charles decides to start helping to initiate other newcomers, just as he himself was initiated. Since he finds it very rewarding, he even signs up both for possible mentorship as well as for knowledge-preaching on the subject of DA-converters, which is the area he has found most interesting, and which he has been focusing on in his own learning process.

5 The ambient web: a glimpse of the future

Figure 8 depicts the five different web layers that we consider most important from a human-computer-interaction perspective. Each layer relies on the ones below it. At the bottom we have the *presentational* (ordinary) web, which is based on the familiar HTML technology. The next layer consists of the *syntactic* (or structural) web, based on XML, which is the dominating standard for structured data exchange between computers. Then there is the (machine) *semantic* web, which is based on RDF, and then– as discussed in Section 3.9 – we have the *human-semantic* (or *conceptual*) web, based on UML.

Figure 8 Five important web layers and their supporting technology

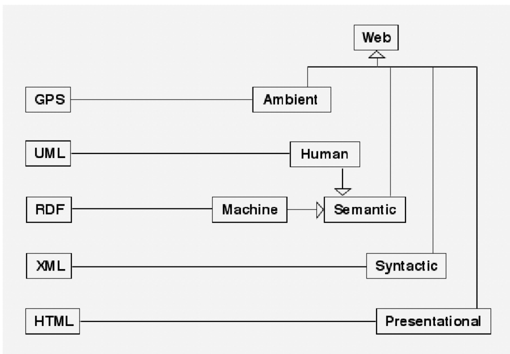

The infrastructure, architecture, frameworks and tools discussed in this paper are concerned with four of the five web layers of Figure 8 – the presentational, the structural, the machine-semantic and the human-semantic (=conceptual) web. However, projects such as *ubiquitous computing* (http://www.ubiq.com/hypertext/weiser/UbiHome.html) and *ambient intelligence* (http://www.eusai.net) imply yet another web layer, which we could call the *ambient web*, by which we mean the personalisable and non-intrusively present web. The possibilities inherent in this web layer will now be briefly discussed.

Let us consider annotating the real world itself by assigning messages to objects and positions. The latter can be achieved by a technique called *mid-air-messaging*[37] which assigns messages to positions by making use of the Global Positioning System. GPS provides unique position identifiers, which are used to create the corresponding URIs. Of course, to assign unique identifiers to physical objects[38] would require some form of agreement on protocol and naming conventions.

Much of this is actually happening in the next-generation technology for mobile phones, PDAs[39] and other types of 'intelligent' objects.[40] Through the use of web technologies, a PDA/mobile phone with built-in GPS and Bluetooth/wireless could locate relevant metadata for a given position and/or object and then display this information to the user. Moreover, this display would not have to be limited to the small screen of the PDA/mobile phone. Instead, one could make use of e.g., a pair of transparent monitor spectacles that can overlay the information in the immediate vicinity of the object or at the correct position.[41] This kind of technology is called *augmented reality* and – to a certain extent – it already exists today.

The metadata for physical objects can also be stored in a personal e-portfolio such as the SCAM portfolio. This makes it possible for you to move around in a physical environment and experience different things depending on which 'perspective' you are assuming – or which mode your 'spectacles' are transmitting to you.[42] Such spectacles could switch between different modes and display information about e.g., the history, the restaurants or the sports events at the place where you are located. They could embed you in fictitious realities and provide interactive 'soap-opera games', or they could expose you to more practically oriented realities that provide useful information about things like road-maps, transportation, biotopes, etc. One of the most obvious uses of this technology is 'personalised spectacles', with which you can annotate your surroundings in support of your own memory. This kind of technology has every potential of becoming a valuable aid for people with memory deficiency problems.

Taken together, the principles and technologies that we have described in this paper form a new paradigm of information flow and access. We have chosen to call this paradigm 'the ambient web'. In the KMR group, we recently have carried out two masters thesis projects that focus on this theme. They concern the splitting of Conzilla into a server-side and a client-side using the J2ME platform, which supports 'thin clients' such as PDAs and mobile phones. This will enable a *mobile conceptual web*, where context-maps on e.g., a PDA can be filled with position-dependent information (content). Of course, the e-learning applications of this type of technology are quite numerous and will not be further discussed here.

6 Conclusions and future work

In this paper we have described our contributions to the infrastructure and the architecture of a Public e-learning platform, which is based on open source and open international ICT standards, and which supports a learner-centric, interest-oriented form of 'knowledge pull'. Moreover, we have shown how our frameworks can help the development of tools that support the knowledge roles of cartographer, librarian and composer, which are three of the seven knowledge roles defined by our educational architecture.

Although we have made considerable progress, much work still remains to be done in order to obtain efficient support for all of the seven knowledge roles. One of the missing pieces is a *workflow engine framework* that could enable the easy construction of customised workflow processes. Another missing component is a *semantic mapping module* as described e.g., by the *SIMILE* project (http://simile.mit.edu/wiki). This could help us transcend the mere 'semantic co-existence' of information that is described in RDF and achieve 'semantic cooperation', i.e., collaboration between systems with different descriptions of the underlying subject domain. A third desirable component is a *semantic matcher*, which would provide support for semi-automatic personalisation of learning material by matching learner profiles with learning object metadata. Important work within this field is being done within the *Prolearn* network (http://www.prolearn-project.org), e.g., by the *ELENA* project (http://www.elena-project.org).

Additions such as these underline the ongoing convergence between the fields of e-learning and knowledge management, which is clearly visible within both research communities, and which is described e.g., in (Lytras et al., 2002). A full-fledged *Public* e-*learning and Knowledge Management Platform* will be of immense value for society, not only for e-learning, but also for e-health, where it will enable patient-centric healthcare, for e-administration, where it will enable 24/7/365 agency services and for e-commerce, where it will enable truly consumer-centric business models.

Acknowledgements

The work of the KMR group is critically depending on several people and organisations that are supporting us in various ways. For organisational and financial support we acknowledge Yngve Sundblad at CID, Barbro Attlestam at VINNOVA, Mia Lindegren at ULL and WGLN, Christina Szekely and Peter Karlberg at MSU, Janiche Opsahl, Agneta Sommansson and Krister Widell at UR, and Allan Westerdahl at CFL. From the research perspective, we acknowledge the contributions to Edutella of L3S in Hanover and AIFB in Karlsruhe. Last but not least we acknowledge the excellent programming contributions to SCAM and SHAME from Jöran Stark, Jan Danils and Henrik Eriksson as well as the profound contributions to Edutella of Pär Sjöberg.

References

Alexander, S. (2001) 'E-learning developments and experiences', *Education + Training*, MCB Univ. Press, Vol. 43, Nos. 4–5, pp.240–248, http://www.emerald-library.com/ft.

Blomqvist, U., Handberg, L. and Naeve, A. (2003) 'New methods for focusing on students learning process and reflection in higher education', *Proc. of the 28th IUT*, June, Improving University Teaching, Conference, Växjö, Sweden.

Decker, S.M.C., Naeve, A., Nejdl, W., Risch, T. and Studer, R. (2001) *Edutella – An Infrastructure for the Exchange of Educational Media*, Part of the PADLR proposal to WGLN, granted in June 2001, www.learninglab.de/pdf/L3S_padlr_17.pdf.

Douglas, J. (1991) *Only Freedom of Education can Solve America's Bureaucratic Crisis of Education*, Cato Policy Analysis No. 155, http://www.cato.org/pubs/pas/pa-155es.html.

Hestenes, D. (1995) 'Modeling software for learning and doing physics', in Bernardini, C. (Ed.): *Thinking Physics for Teaching*, pp.25–65, http://modelingnts.la.asu.edu/pdf/ModelingSoftware.pdf.

Hestenes, D. (2002) *Reforming the Mathematical Language of Physics*, Oersted Medal Lecture 2002 – in response to his reception of the Oersted medal award, http://modelingnts.la.asu.edu/pdf/OerstedMedalLecture.pdf.

Knudsen, C. and Naeve, A. (2001) 'Presence production in a distributed shared virtual environment for exploring mathematics', *Proc. of the 8th Intl. Conf. on Advanced Computer Systems (ACS 2001)*, Szcecin, Poland, October 17–19, http://kmr.nada.kth.se/papers/TelePresence/CJKPresenceProd.pdf.

Kraan, W. (2003) *Using SHAME to Fill your SCAM*, Center for Educational Technology Interoperability Standards, November, http://www.cetis.ac.uk/content2/20031105152216.

League of World Universities (1993) 'Large city universities and the public school system', *Rectors Conference*, New York University, April 19–20, www.nyu.edu/rectors/conf41993A.html.

Lytras, M., Pouloudi, A. and Poulymenakou, A. (2002) 'Knowledge management convergence: expanding learning frontiers', *Journal of Knowledge Management*, Vol. 6, No. 1, pp.40–51.

Naeve, A. (1997) *The Garden of Knowledge as a Knowledge Manifold – A Conceptual Framework for Computer Supported Subjective Education*, CID-17, TRITA-NA-D9708, CID/NADA/KTH, Stockholm, http://kmr.nada.kth.se/papers/KnowledgeManifolds/cid_17.pdf.

Naeve, A. (1999) *Conceptual Navigation and Multiple Scale Narration in a Knowledge Manifold*, CID-52, TRITA-NA-D9910, CID/NADA/KTH, Stockholm, http://kmr.nada.kth.se/papers/ConceptualBrowsing/cid_52.pdf.

Naeve, A. (2001a) 'The concept browser – a new form of knowledge management tool', *Proc. of the 2nd European Web-Based Learning Environment (WBLE) Conference*, Lund, Sweden, October 24–26, pp.151–161, http://kmr.nada.kth.se/papers/ConceptualBrowsing/ConceptBrowser.pdf.

Naeve, A. (2001b) 'The knowledge manifold – an educational architecture that supports inquiry-based customizable forms of e-learning', *Proc. of the 2nd European Web-Based Learning Environments (WBLE) Conference*, Lund, Sweden, October 24–26, pp.200–212, http://kmr.nada.kth.se/papers/KnowledgeManifolds/KnowledgeManifold.pdf.

Naeve, A. and Nilsson, M. (2004) 'ICT-enhanced mathematics education within the framework of a knowledge manifold', *Proc. of the 10th International Congress of Mathematics Education (ICME)*, Copenhagen, Denmark, July 4–11, http://kmr.nada.kth.se/papers/MathematicsEducation/ICME2004-ICT-enhanced-math-ed.pdf.

Naeve, A. and Selg, H. (2001) *Dynamisk Handbok för Statens Energimyndighet*, KTH, Stockholm, http://kmr.nada.kth.se/papers/ConceptualBrowsing/KMR-myndighetshandbok.pdf.

Naeve, A., Nilsson, M. and Palmér, M. (2001) 'The conceptual web – our research vision', *Proc. of the First Semantic Web Working Symposium (SWWS'01)*, Stanford University, California, USA, July 30–August 1, pp.72–73, www.semanticweb.org/SWWS/program/position/soi-nilsson.pdf.

Nejdl, W., Wolf, B., Qu, C., Decker, S., Sintek, M., Naeve, A., Nilsson, M., Palmér, M. and Risch, T. (2002) 'Edutella: A P2P networking infrastructure based on RDF', *Proc. of the 11th World Wide Web Conference (WWW2002)*, Hawaii, USA, May 7–11, http://kmr.nada.kth.se/papers/SemanticWeb/p597-nejdl.pdf.

Nilsson, M. (2000) *The Conzilla Design – The Definitive Reference*, CID/NADA/KTH, Stockholm, http://kmr.nada.kth.se/papers/ConceptualBrowsing/conzilla-design.pdf.

Nilsson, M. (2001) *The Semantic Web: How RDF will Change Learning Technology Standards*, Feature article, Centre for Educational Technology Interoperability Standards (CETIS), October, http://kmr.nada.kth.se/papers/SemanticWeb/CETISSemWeb.pdf.

Nilsson, M., Palmér, M. and Braase, J. (2003) 'The LOM-RDF binding – principles and implementations', *Proc. of the 3rd Annual Ariadne Conference*, Katholieke Universiteit Leuven, Belgium, November 20–21, http://kmr.nada.kth.se/papers/SemanticWeb/LOMRDFBinding-ARIADNE.pdf.

Nilsson, M., Palmér, M. and Naeve, A. (2002) 'Semantic web metadata for e-learning – some architectural guidelines', *Proc. of the 11th World Wide Web Conference (WWW2002)*, Hawaii, USA, May 7–11, http://kmr.nada.kth.se/papers/SemanticWeb/p744-nilsson.pdf.

Nilsson, M., Naeve, A. and Palmér, M. (2004) 'The edutella P2P network – supporting democratic E-Learning and communities of practice', in McGreal, R. (Ed.): *Online Education Using Learning Objects ISBN 0-415-33512-4*, Routledge-Falmer, New York, http://kmr.nada.kth.se/papers/SemanticWeb/Edutella-chapter.pdf.

Palmér, M., Naeve, A. and Nilsson, M. (2001) 'E-learning in the semantic age', *Proceedings of the 2nd European Web-Based Learning Environment (WBLE) Conference*, Lund, Sweden, October 24–26, pp.162–174, http://kmr.nada.kth.se/papers/SemanticWeb/e-Learning-in-The-SA.pdf.

Palmér, M., Naeve, A. and Paulsson, F. (2004) 'The SCAM framework – helping applications to store and access metadata on the semantic web', *Proceedings of the First European Semantic Web Symposium*, ISBN 3-540-21999-4, Springer, Heraklion, Greece, May 10–12, http://kmr.nada.kth.se/papers/SemanticWeb/SCAM-ESWS.pdf.

Paulsson, F. and Naeve, A. (2003) 'Standardized content archive management – SCAM', *IEEE Learning Technology Newsletter*, Vol. 5, No. 1, pp.40–42, http://lttf.ieee.org/learn_tech/issues/january2003/learn_tech_january2003.pdf.

Pettersson, D. (2000) *Aspect Filtering as a Tool to Support Conceptual Exploration and Presentation*, TRITA-NA-E0079, CID/NADA/KTH, Stockholm, http://kmr.nada.kth.se/papers/ConceptualBrowsing/AspectFiltering-exjobb.pdf.

Powell, A., Nilsson, M., Naeve, A. and Johnston, P. (2004) *DCMI – Abstract Model (Working Draft)*, http://www.ukoln.ac.uk/metadata/dcmi/abstract-model.

Rumbaugh, J., Jacobsson, I. and Booch, G. (1999) *The Unified Modeling Language Reference Manual*, Addison Wesley Longman Inc., ISBN 0-201-30998-X.

Wilson, S. (2001) *The Next Wave: CETIS Interviews Mikael Nilsson about the Edutella Project*, Feature article, Centre for Educational Technology Interoperability Standards (CETIS), September, http://kmr.nada.kth.se/papers/SemanticWeb/TheNextWave.pdf.

Notes

[1]This presentation is condensed from (Hestenes, 2002).

[2]Hestenes attaches the following warning to this (constructivist) learning principle:

> "There are many brands of constructivism, differing in the theoretical context afforded to the constructivist principle. An extreme brand called 'radical constructivism' asserts that constructed knowledge is peculiar to an individual's experience, so it denies the possibility of objective knowledge. This has radicalised the constructivist revolution in many circles and drawn severe criticism from scientists. I see the crux of the issue in the fact that the constructivist principle does not specify *how* knowledge is constructed. When this gap is closed with the other learning principles and scientific standards for evidence and inference, we have a brand that I call *scientific constructivism*".

[3]In KM terminology, we use the term 'knowledge component', when we assume the perspective of the teacher(s), the term 'information component' (or 'learning component') when we assume the perspective of the learner(s), and the term 'resource component' when want to remain neutral in this respect. It is the *transformation* of learner information into learner knowledge, and eventually the *transmutation* of learner knowledge into learner understanding, that is the overall aim of the learning process.

[4]The term 'metadata ecosystem' is explained in Section 3.9.

[5]This supports the separation of content from context, which promotes the reuse of content across different contexts.

[6]Note that when maintaining his or her knowledge patch, the gardener performs the role of cartographer when working with the contexts and the role of librarian when working with the content.

[7]The term 'semantic' is roughly synonymous with 'understandable'.

[8]The parenthesised words in this list refer to the KM knowledge roles presented in Section 2.3.

[9]Within the semantic web community usually called *ontologies*.

[10]World Wide Web Consortium.

[11]i.e., has machine-processable semantics.

[12]A thorough discussion of this misconception of metadata is given in Nilsson *et al*. (2002).

[13]Personalised access to distributed learning repositories.

[14]That coordinates the PADLR project.

[15]Where the KMR group is participating together with all of the mentioned research groups and institutes.

[16]Often referred to as 'e-learning silos'.

[17]With differences in up-time, performance, storage size, functionality, number of users, etc.

[18]Note that there are several occurrences of 'or' in this transcription. However, this information is not explicit in the figure, but is represented separately.

[19]Developed by e.g., Dublin Core (http://dublincore.org), IMS (http://www.imsproject.org), IEEE/LTSC (http://ltsc.ieee.org), and ISO/IEC-JTC1/SC36 (http://jtc1sc36.org).

[20]Universal Resource Identifiers.

[21]These RDF-bindings where developed under the coordination of Mikael Nilsson as part of the KMR involvement in metadata standardisation work within Dublin Core (Powell *et al*., 2004), IMS (http://www.imsglobal.org/metadata), IEEE (http://ltsc.ieee.org/wg12), and ISO (http://jct1sc36.org). Our ongoing work within this field is described in our draft of the IEEE LOM-RDF binding (http://kmr.nada.kth.se/el/ims/metadata.html).

[22]This specification was developed by Mikael Nilsson and Matthias Palmér in order to solve this problem (and others) for the SCAM project as well as for the Edutella project.

[23]Such as the IMSEVimse LOM-editor (http://kmr.nada.kth.se/imsevimse).

[24]If there are any restrictions.

[25]As long as the metadata is expressed in RDF.

[26]VWE is being developed by the department for Interactive Media and Learning (IML) at Umeå University in collaboration with the European School net (EUN). The project is lead by Fredrik Paulsson.

[27]See Section 2.3.

[28]The ease with which such filters are constructed and modified is a major strength of the Conzilla tool.

[29]About 300 of these learning objects have been created by Ambjörn Naeve. They are also available through his Mathemagic™ Edufolio (http://knowgate.nada.kth.se:8080/portfolio/main?manifest=amb&cmd=open), which is freely licensed for non-commercial use under a Creative Commons license (http://creativecommons.org).

[30]For an explanation of this ULM notation, see Figure 1.

[31]Certified identities will be crucial in order to establish the 'web of trust' that is a pre-requisite for realising the full potential of e-commerce, and which is just as vitally important for effective e-learning.

[32]See Section 2.3.

[33]Such a system is presently under construction at CID. An interesting experiment in this direction, making use of advanced technology for presence production, is described in (Knudsen and Naeve, 2001).

[34]Empirical studies such as Alexander (2001) have found that personal feedback is one of the features of e-learning that students appreciate the most.

[35]Conzilla 1 has been developed by Mikael Nilsson and Matthias Palmér, and Conzilla 2 is being developed by Matthias Palmér and Henrik Eriksson.

[36]We are also aiming for Conzilla to support increased e-accessibility, in accordance with the e-Europe initiative (http://europa.eu.int/information_society/eeurope/2005/text_en.htm), by enabling the program to adapt itself to different cognitive profiles.

[37]Developed by Hewlett-Packard Bristol Labs (http://www.interex.org/hpworldnews/hpw203/03news.html).

[38]Independently of their position.

[39]Personal Digital Assistant.

[40]E.g., refrigerators with their 'own' websites, etc.

[41]From the perspective of the observer.

[42]i.e., which metadata sources that you are including/filtering.

[43]All listed URLs have been accessed on 11th July 2004

List of abbreviations

AIFB	Institut für Angewandte Informatik und Formale Beschreibungsverfahren
API	Application Programming Interface
CFL	Swedish National Centre for Flexible Learning
CID	Centre for user-oriented Information Technology Design
DCMI	Dublin Core Metadata Initiative
ECIMF	Electronic Commerce Integration Meta Framework
GPL	GNU Public License
GNU	Gnu is Not Unix
GPS	Global Positioning System
GUI	Graphical User Interface
ICT	Information and Communication Technology
IEEE	Institute of Electrical and Electronics Engineers
IMS	Instructional Management Systems
JSP	Java Server Pages
JXTA	JuXTApose
J2EE	Java 2 Enterprise Edition
J2ME	Java 2 Micro Edition
HTML	HyperText Markup Language
HTTP	HyperText Transfer Protocol
KM	Knowledge Manifold
KMR	Knowledge Management Research
KTH	Royal Institute of Technology
LGPL	Light GNU Public License

LMS	Learning Management System
LOM	Learning Object Metadata
LTSC	Learning Technology Standards Committee
L3S	Learning Lab Lower Saxony
MPL	Mozilla Public License
MSU	Swedish National Agency for School Improvement
OAI	Open Archives Initiative
OLR	Open Learning Repository
PADLR	Personalised Access to Distributed Learning Resources
PDA	Personal Digital Assistant
PeLP	Public e-Learning Platform
P2P	Peer-to-Peer
RDF	Resource Description Framework
RDFS	Resource Description Framework Schema
SCAM	Standardised Content Archive Management
SCORM	Sharable Content Object Reference Model
SHAME	Standardised Hyper-Adaptable Metadata Editor
SOAP	Simple Object Access Protocol
ULL	Uppsala Learning Lab
ULM	Unified Language Modelling
UML	Unified Modelling Language
UR	Swedish Educational Broadcasting Company
URI	Universal Resource Identifier
URL	Universal Resource Locator
URN	Universal Resource Name
VINNOVA	Swedish Agency for Innovation Systems
VML	Virtual Mathematics Laboratorium
VWE	Virtual Workspace Environment
WGLN	Wallenberg Global Learning Network
W3C	World Wide Web Consortium
XML	eXtensible Markup Language

Websites

AIFB (at Universität Karlsruhe), http://www.aifb.uni-karlsruhe.de/english.

Ambient Intelligence, http://www.eusai.net.

Ambjörn's Mathemagic™ ConFolio, http://knowgate.nada.kth.se:8080/portfolio/main?manifest=amb&cmd=open.

CID (Centre for user oriented IT Design), http://cid.nada.kth.se/en.

Concept maps, http://www.graphic.org/concept.html.

Conzilla, www.conzilla.org.

CourseWare Watchdog (AIFB Karlsruhe), http://www.aifb.uni-karlsruhe.de/Publikationen/showPublikationen?id_db=51.

Creative Commons, http://creativecommons.org.
Digital Media Library (UR), http://www.ur.se/mb.
Dublin Core Metadata Initiative (DCMI), http://dublincore.org.
ECIMF, http://www.ecimf.org.
EducaNext, http://www.educanext.org.
Edutella, http://edutella.jxta.org.
e-Europe, http://europa.eu.int/information_society/eeurope/2005/text_en.htm.
ELENA, http://www.elena-project.org.
Hestenes' modelling website, http://modelingnts.la.asu.edu/html/Modeling.html.
IEEE/LTSC, http://ltsc.ieee.org.
IEEE-LOM, http://ltsc.ieee.org/wg12.
IMS Content Packaging, http://imsglobal.org/content/packaging.
IMS Metadata, http://www.imsglobal.org/metadata.
IMS, http://www.imsproject.org.
IMSEVimse (LOM-editor), http://kmr.nada.kth.se/imsevimse.
ISO/IEC JTC1 SC36, http://jtc1sc36.org.
JXTA, http://www.jxta.org.
KMR (Knowledge Management Research) group, http://kmr.nada.kth.se.
L3S (Learning Lab Lower Saxony) Research Center, http://www.learninglab.de.
Learn in Freedom, Books Critiquing the School System, http://learninfreedom.org/system.html.
Learning Resource Centre (CFL), http://larresurs.cfl.se.
IEEE LOM-RDF-binding, http://kmr.nada.kth.se/el/ims/metadata.html.
Mid-Air-Messaging, http://www.interex.org/hpworldnews/hpw203/03news.html.
PADLR, http://www.learninglab.de/padlr/index.html.
Prolearn, http://www.prolearn-project.org.
RDF (Resource Description Framework), http://www.w3.org/RDF.
SCAM, http://scam.sourceforge.net.
Semantic web, http://www.semanticweb.org.
SHAME, http://kmr.nada.kth.se/shame.
SOAP, http://www.w3.org/TR/SOAP.
Soft Infrastructure for IT in education (MSU), http://mjukis.skolutveckling.se.
SWEBOK consumer, http://edutella.jxta.org/downloads.
The SIMILE project, http://simile.mit.edu/wiki.
Topicmaps, http://www.topicmaps.org.
Ubiquitous Computing, http://www.ubiq.com/hypertext/weiser/UbiHome.html.
UDBL (Uppsala DataBase Laboratory), http://user.it.uu.se/~udbl.
ULL (Uppsala Learning Lab), www.ull.uu.se.
UML (Unified Modelling Language), http://www.uml.org.
UNIVERSAL, http://www.ist-universal.org.
VWE (Virtual Workspace Environment), www.vwe.nu.
Web Services, http://www.w3.org/2002/ws.
WGLN (Wallenberg Global Learning Network), www.wgln.org.