Technologies for Semantic Project-Driven Work Environments

Bernhard Schandl
Department of Distributed and Multimedia Systems
University of Vienna, Austria
bernhard.schandl@univie.ac.at

Ross King, Niko Popitsch
Austrian Research Centers GmbH – ARC Research Studios
Studio Digital Memory Engineering, Vienna, Austria
ross.king | niko.popitsch@researchstudio.at

Brigitte Rauter, Martin Povazay
P.Solutions Informationstechnologie GmbH, Vienna, Austria
brigitte.rauter | martin.povazay@psolutions.at

Keywords: knowledge capture, data integration, collaborative semantics

1. Introduction

As computer and Internet applications became ubiquitous, most daily business must handle an increasing amount of information via several applications and systems such as e-mail applications, file systems, business software, databases, or other systems. Dealing with information flows is not restricted to a special skill level or field of work; it is rather a significant attribute of any computational work environment.

Most ongoing tasks in companies are in the context of a project since management strategies force process-driven business and organization. During a project a large volume of knowledge arises that is connected to the output of the project (products and services) as well as to know-how regarding the project realisation and the use of resources.

The important aspect of capturing this knowledge in some form is to impart and recycle organizational knowledge to gain raising efficiency in doing business. This knowledge capture may result in the creation of (digital) documents like e.g. project plans, resource plans, reports, product sheets, and so on. But there is also a second aspect of knowledge capture, regarding a kind of *semantic glue* between those assets that should be captured in order to be able to relate them to each other (e.g. who did what on a certain project, how do requirement documents and project reports relate, ...), that is, to define their context.

The present…

This second aspect of knowledge capture is ignored by most systems and methodologies in place today. Consider for example the common practice of storing the majority of project documents on a shared file server. Semantic information about the *meaning* of a certain document or about its relation to other documents in the same or in other projects can only be captured in a most restricted way (e.g. by using file name and/or file path conventions or by describing such relations in these documents themselves).

The result of this approach is that it is very cumbersome to find documents on such a file server as soon as it grows to a certain size. This document management strategy supports the finding of documents only by browsing a strictly hierarchical directory structure that follows a certain naming convention, or by searching for low-level metadata features (e.g. creation date) or for some text in the document (full text search).

Specialized software (e.g. project management software) was meant to overcome these shortcomings. Hundreds of tools have been developed in this area – nevertheless most of them are fairly closed-box systems that are difficult to customize and that force a company to shape its business processes to fit the software's requirements, rather than the other way around. Furthermore such systems—where used—do not fulfil all of the requirements in a project-driven working environment.

The future…

It is our opinion that highly collaborative semantic systems are needed to advance the state of the art in knowledge capture and reuse in the context of project and document management. Any system that supports users in managing their documents must be able to *capture the semantic glue* between these documents in some form. We suggest that a mixture between natural language (for users) and formal description languages (for computers) would be very beneficial here.

Capturing the semantics of documents and their interrelations supports *finding*, *exploring*, *reusing* and *exchanging* digital documents. Furthermore, this context information may be an essential aspect of the long-term preservation of such these documents. We believe that the process of capturing semantics must take place when the system users have maximum knowledge about a certain document (i.e. when the document is created or updated) and should interfere with a user's normal workflow as little as possible. As every organization has slightly different internal workflows and requirements for such a system, the underlying software must be highly configurable and easily adapted towards the organization's (changing) needs. Furthermore—as project work is always team work—we want to emphasize the need for a strong collaborative character of such systems.

The various aspects of information capture and the distributed nature of the utilized information sources furthermore demands an infrastructure that supports interconnection and integration of multiple heterogeneous data sources. We observe a demand for semantic systems in areas in which knowledge work and collaboration is required, for example for managing liability cases, audit reports, or inspection reports; in software development, product management, management consultancy, or innovation management. Knowledge is one of the most important assets of organisations in these fields, which accounts for the demand for semantic (knowledge) work environments.

In this chapter, we will illustrate how a suite of semantic technologies can help to decrease the effort required for knowledge organization, storage, and retrieval.

2. Working in Project-Driven Environments: The Need for Collaborative Semantics

Before delving into technical details, we will give a short description of typical practices and processes in project-driven environments in industry in order to show where we see potential for capturing information in its semantic context. This example will demonstrate the apparent need for collaborative semantics.

In general, a project is a compilation of specific tasks and activities that are required in order to reach a defined goal within a defined time frame. A project is usually realized by a team of people with different skills and experience. Project-oriented work environments face the need of dealing with a very large volume of data stored in various files or system applications: material bills, object lists, detailed descriptions and specifications, and other resources. When working in a team there is increased demand for communication between team members and, in particular, demand for the documentation of communication.

Following our own project experience during the last years, we can (amongst others) identify the following questions that frequently arise when working in small- to medium-sized IT projects:

- What kind of project information is stored in which system, and what meaning has such information in which context?
- What are the naming and structure conventions for information, e.g. in a file system?
- Who is responsible for which information and task, and which time constraints have to be considered?
- What kind of dependencies between information and tasks are observable?

Everyday work consists, to a certain extent, of repeating *mini projects* or processes where, in our opinion, semantically enriched information can be captured automatically by a semantic management system, assuming that such a system does not overburden the project team with additional effort. The tracking of information in semantic contexts starts with the definition of the project aim and the planning of the realisation. Typically, this information is stored in file systems, project management systems or databases. Consequently, the basic ontology for a project is already defined during the project's initial stage, including all persons that are assigned to tasks based on their skills, experiences and responsibilities.

In order to reach the defined milestones, the team members focus on how to realize the given tasks. They must analyze what is new and what can be reused. The project will be itemized: Which components are needed to develop an application with its specific attributes, depending on its usage context. They look after successfully realized projects with similar problems: Is this knowledge already available? Were there any similar situations? Who was involved and what were the detailed specifications of material and costs?

In the next stage, the project sub-teams collect and develop specific information to reach their goals, therefore extending and refining the project ontology. During all process steps different systems are used to represent information and semantics: groupware systems like Outlook, Exchange, ERP systems, CRM systems; project management tools like MS Project, intranet services, file systems, and others. Often, there is no entity that integrates information from these various sources and presents a unified view on this information to the user.

In this chapter, we describe a set of technologies that provide an infrastructure for the realization of small- and medium-sized projects. We describe METIS, a database that can act as integration point for ontologies and data gathered from external systems, and Ylvi, a semantic wiki that allows creation of semantic information during normal project work. The SemDAV protocol, introduced in this chapter, enables communication and the exchange of data and metadata between subsystems. The Semplorer, a graphical user interface constructed on top of this protocol, provides different views and management functions on the unified set of project-related information. In the following, we illustrate how these technologies can be used together to cope with the efforts of project-driven work environments.

3. Related Work

The integration of information from different sources in order to create a unified organizational memory has been studied in various works

(Wache, Vögele, Visser, Stuckenschmidt, Schuster, Neumann & Hübner, 2001; Noy, 2004). Semantic Web technology provides a suitable basis for such a task (Priebe & Pernul, 2003), since it provides a data model for representing nearly every kind of information (RDF[1]), both theoretical and practical models that allow further processing of metadata (namely ontological organizing and reasoning), and implementations of these models that are being developed by an increasingly large community.

Ontology-based Information Integration

Examples for ontology-based information integration systems are Observer (Mena, Illarramendi, Kshyap & Sheth, 2000) or the RDF-based Ontobroker (Decker, Erdmann, Fensel & Studer, 1998). The central approach of these systems is the integration of heterogeneous sources by making the semantics of data explicit through the definition of ontologies. Their goal is to enable semantic (or concept-based) queries over the integrated data sources via a uniform query interface. Structural and semantic conflicts between the ontologies that are involved in integration scenarios can are reconciled through the definition of mappings. An important representative for mapping frameworks built on Semantic Web technologies is the MAFRA framework (Maedche, Motik, Silva & Volz, 2002).

Recent works on information integration focus on the requirements of organisations consisting of many scattered units. Peer-to-peer (P2P) networks are a possible approach to fulfil these requirements and systems like Edutella (Nejdl, Wolf, Changtao, Decker, Sintek et al., 2002) or GridVine (Aberer, Cudre-Maurox, Hauswirth & van Pelt, 2004) are knowledge based integration systems for P2P architectures. The EU Project BRICKS (Risse, et al, 2005) combines this peer-to-peer approach with the semantic integration techniques mentioned above,

Integration of information requires access to the data stored in organizational data sources. D2RQ (Bizer, Seaborne, 2004) is a framework for treating data in relational databases as RDF graphs and Lethi (Lethi, Frankhauser, 2004) propose an approach for XML data integration with OWL.

An example of the useful integration of semantic web technology and web services is the SemanticLIFE project (Anjomshoaa, Manh Nguyen, Shayeganfar, & Tjoa, 2006). In this work, data from different sources,

[1] RDF: `http://www.w3.org/RDF/`

including Outlook databases or instant messaging clients, is integrated with the aim to build a semantic repository of a human's personal information. We believe that this idea can be applied to organizational knowledge as well, keeping in mind that it is still important to consider an individual's personal information sphere. We hypothesize that a significant portion of organizational knowledge is concealed in seemingly personal information like e-mail messages or instant messaging communication, and every approach for organization-wide knowledge management must consider these types of information.

The Semantic Desktop

The above hypothesis is supported by the increasing importance of research that aims at the creation of the semantic desktop (Sauermann, Bernardi, & Dengel, 2005), a new metaphor for human-centered information management using semantic technology. Important projects in this field are Haystack (Karger, Bakshi, Huynh, Quan, & Sinha, 2005), Gnowsis (Sauermann, 2005), Chandler (Fitzgerald, 2003), and DeepaMehta. The common denominator of most semantic desktop projects is the use of ontologies for data organization and the storage of metadata in RDF, which makes the systems open for interconnection to other platforms, and allows the composition of organization-wide semantic systems. Another possible approach for data integration from heterogeneous systems is described by Haslhofer (2006), where SPARQL queries and corresponding result graphs are rewritten on-the-fly by mediators, according to predefined schema mappings.

Organizational intranets are often considered a primary storage pool for organizational knowledge. Recent studies (Géczy, Izumi, Akaho, & Hasida, 2006; Lamb, Davidson, 2005) indicate that knowledge workers often make use of only a limited subset of intranet resources. By applying semantic web technologies not only to the World Wide Web and the desktop, but also to intranet resources, it will be possible to create a technology bridge between these different worlds. As there will be no single point of integration in such environments, the peer-to-peer paradigm may serve as underlying networking paradigm. Peer-to-peer technology for knowledge management has been presented e.g. by Kaulgud & Dolas (2006) and Le Coche, Mastroianni, Pirrò, Ruffolo, & Talia (2006).

(Semantic) Wikis

Another approach to facilitate semantic knowledge exchange within organizations and/or communities of practice is the usage of wikis (Ajchrzak, Wagner, & Yates, 2006) or semantic wikis (Tolksdorf & Simperl, 2006; Oren, Breslin, & Decker, 2006; Völkel, Krötzsch, Vrandecic, Haller, & Studer, 2006). In order to find a common agreement on how to interchange data between different semantic wiki systems, Völkel & Oren (2006) proposed the Wiki Interchange Format (WIF). However, many wiki implementations suffer from two principle problems:

(1) With the increasing number of (semantic) features within wikis, the user interface (i.e. the wiki markup) becomes more complicated, which discourages the use of these features. Also, semantic features are often poorly documented, and the semantics of using them is not clearly defined. Intuitive and clear user interfaces are required in order to encourage users and further promote the idea of semantic wikis.

(2) Many wiki implementations lack support for management of data that cannot be represented by text-based formats, e.g. multimedia data, and integration of legacy systems (e.g. database-driven CRM systems) is difficult.

We consider (semantic) wikis as a valuable tool for information organization and collaborative knowledge work. In order to increase user acceptance, the abovementioned problems still have to be investigated.

4. Technologies for Engineering an Organizational Memory

In this section, we present an overview of the architecture and relevant technical details of: METIS, a framework for the management of multimedia data and metadata; Ylvi, a semantic wiki platform built on top of METIS; and SemDAV, a Semantic-Web-based protocol that allows the integration of various data sources into a unified information model. We illustrate use cases for these technologies in the context of project-driven work environments and point out how these systems may work together to give enterprise-wide tool support for knowledge tasks.

Architecture Overview

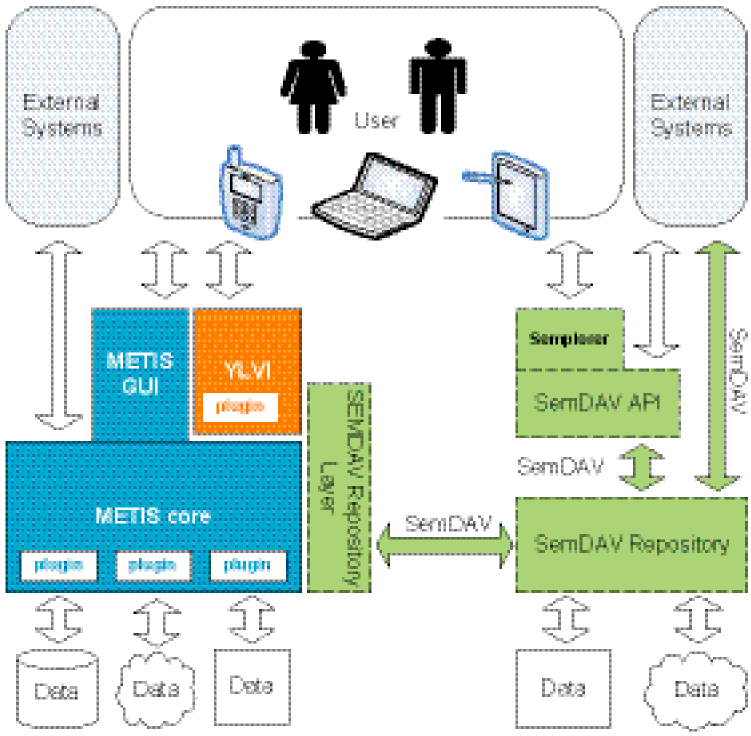

Figure 1: Knowledge Support for Project-Driven Work Environments

In Figure 1, we give an overview of our envisioned architecture for project-driven work environments. The semantic database for multimedia objects METIS provides support for persistence of data and metadata as well as a flexible plugin framework for content analysis. On top of METIS, a web-based Wiki-like interface (Ylvi) provides user access to stored data and annotations through configurable rendering pipelines. Existing desktop applications may be integrated into a semantic repository where enterprise data is represented using Semantic Web technology (SemDAV) which may also be backed by a METIS instance. SemDAV provides a protocol and an Application Programming Interface so that it can be used as semantic storage system for any business application. Through the Semplorer, a graphical user interface for

SemDAV, users may access data using similar interaction metaphors like common file browsers provide.

In the next sections, we introduce the individual components of our system suite in more detail and illustrate their usage with the continuous example of small- and medium project work. Despite the fact that some aspects of our architecture are still under development, we have already carried out research and development projects using our system suite, and initial experience showed that projects benefit from the slim, efficient structures fostered by our tool suite.

METIS: A Flexible Database Foundation for Unified Media Management

Currently digital data—both in the internet and in the intranet domain—is shifting more and more towards multimedia data. Reasons for this include

- widely spread media capturing technologies (e.g. digital cameras, video enabled cell phones, ...)
- affordable mass-storage and storage devices
- increasing bandwidth for transportation of digital content

It is becoming more and more unusual to store data purely in text formats. The increased availability of digital multimedia content demands powerful data management and processing technologies that take the special properties of this kind of data into account.

Multimedia management is different from traditional text-based methods. Many techniques that are easily applied to text (e.g. indexing, comparison, transformation) are hardly applicable to other media formats or require quite different approaches (for example image comparison or video indexing). Most current approaches rely on the capturing and appliance of metadata of some form to fulfil these tasks. This requires infrastructure that is able to

- capture the data and its metadata in appropriate formats,
- apply domain and media specific algorithms on the managed data and metadata, and

- expose the managed data to external systems/applications in well-defined exchange formats.

> **Example** In project work, a significant fraction of information in a common working environment is hidden in email attachments – some of the attached media items' formats are indexable by the email client, but many are not; thus, searching for email attachments is still cumbersome. Furthermore, the implicit and explicit relations between various media items can be used to support document discovery; for example, one could navigate from a Microsoft Word document on a local hard disk to the email message(s) this document was attached to, and to other documents by the same sender, or to other related emails and attached documents.

As a first step towards addressing the mentioned problems with current state-of-the-art multimedia management technologies, we have developed a prototypical infrastructure that supports uniform and semantic multimedia management for a broad range of possible multimedia applications. Typical multimedia database systems focus on specific kinds of media and/or applications (e.g. video databases, image databases). In the following we provide a short overview of METIS, a database foundation for the unified management of multimedia data of all kinds, comprehensively documented by King, Popitsch, & Westermann (2004, 2007).

The major strength of METIS is its flexibility in adapting to a large number of possible multimedia application scenarios. This is achieved by providing highly flexible frameworks for customizing and extending the implemented semantic data model as well as the associated functionality (including metadata extractors and comparison functions; but also interfaces to external systems).

Semantic data model. METIS is based on an expressive data model that can instantiate any desired scheme for media management, description, and classification. The instantiated data model can be extended/changed at any time. So-called semantic packs permit the bundling of domain-specific customizations and the introduction of widely-used metadata standards (e.g. Dublin Core (ISO 2003), or MARC). Semantic packs can be created through an internal build environment or directly exported from Protégé (Noy, Sintek, Decker, Crubezy, Fergerson, & Musen, 2001).

METIS includes a query processor that supports hybrid searches (Wen, Li, Ma, & Zhang, 2003) for media objects, taking into account the semantic classification of media, their high-level characteristics and low-level features, and their relationships to other media objects.

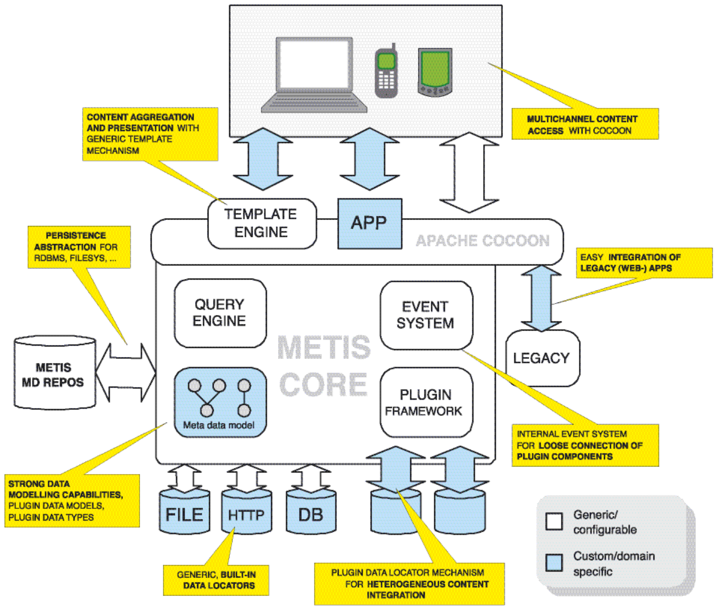

Figure 2: METIS Architecture Overview

Persistence abstraction. METIS is built on a persistence abstraction layer that facilitates the switching of storage back-ends and offers a customizable web front-end for administration and media management. An extensible locator mechanism allows for transparent access to media in different storage systems and locations.

Media aggregation and production. On the output side, METIS offers an XSLT-based multi-channel publishing strategy that can be used for aggregation of the managed media objects and delivery in various context-dependent output formats, like XHTML, SMIL, or Microsoft Word. This production layer can easily be extended to deliver content in arbitrary formats (e.g. RDF) to external applications.

Plugins. A sophisticated plugin infrastructure enables extensive customizations of the system core, such as media- and domain-specific query operators, similarity measures, and feature extraction algorithms. The implemented frameworks provide plugin interfaces for

- complex data types (e.g. MPEG-7 media descriptors (Martínez, Koenen, & Pereira, 2002))
- data model extensions
- querying of the contained data
- metadata and data comparison
- transparent media locators (see above)
- media aggregation and presentations (see above)
- event based communication
- visualization and persistence abstraction (see above)

METIS differs clearly from traditional media-type specific databases, including video or image databases as these are tailored towards their respective media type but do not allow uniform multimedia management. METIS differs from more comparable systems mainly in its flexibility and extensibility on all architectural levels (back-end, GUI, data model and functionality) that greatly helps METIS to quickly fit into concrete projects.

Ylvi: A Semantic Multimedia Wiki Framework

In the previous section, we described how the system/platform METIS can provide uniform access to media resources by providing transparent mechanisms to locate resources in arbitrary data sources, semantic annotation and interrelation of these resources, and querying and aggregated delivery of this media. Such an infrastructure may be sufficient for machines, but human users require additional features, including navigation between media items without the need for searching, or human-readable annotations of media and relations. The *semantic glue* between media objects is stored *in the mind of the users*. The question arises: How can this glue be made explicit without penalising the human in favour of the machine?

> **Example** Typical project documentation consists of multiple documents and media objects that serve completely different

documentation purposes (e.g. a project plan, a resource plan, deliverables, material, emails). Especially when the project's size renders it unprofitable to establish dedicated organization structures, these documents are stored in a common project directory tree, a database, and/or an email server, respectively. Furthermore, these items are stored in various media formats that require the underlying management infrastructure to be able to deal with these formats.

As mentioned before, this representation of project data has various drawbacks, as it does not represent any relation between these documents (aside from that they belong to the same project and are therefore stored in the same folder – if we consider only the project folder). This makes it impossible to search or navigate along such relations or e.g. request a list of all managed documents of a certain type. Data integration platforms like METIS can provide uniform access to such resources and allow the users to model the relationships and semantic properties (e.g. what purpose a certain document serves) that describe them. This enables users to search and exchange this information, but the following questions remain:

- What kind of user interface should be used to enter/access this information?
- What information should be used to browse the data?
- How to describe the semantic glue between data objects in a human-readable fashion?

Although, for example, specialized project management software could provide such functionality, such systems are in general very rigid and tailored towards their respective tasks. They provide only limited interoperability with other systems (e.g. email systems, intranet applications, or 3rd party document management software), are often difficult to extend, and lack flexibility.

Most relations between media objects can be explicitly modelled with systems like METIS (or the mentioned specialized software with all described drawbacks). However, it is not always preferable to represent such relations only with a formal data model. Although machines can work with this information, it is hard for human users to understand such formal relationships, even if the system tries to translate them into human readable representations. It is often more desirable to allow humans to annotate media and its interrelations in natural language text and additionally allow them to model the formal media relations in a simple manner – this is the goal of semantic wikis.

Ylvi is a semantic wiki framework based on the METIS platform. It combines the advantages of structured, type-/attribute-based media management and the open, relatively unstructured wiki approach. By representing wiki pages as abstract media objects, Ylvi can offer sophisticated media management features to the wiki domain and provide an extensible, highly configurable and multimedia-enabled semantic wiki.

The manifold extensibility of the underlying METIS framework renders Ylvi more a *semantic wiki toolkit* than just another semantic wiki instance. The semantic typing of Ylvi media, articles and links between such items greatly enhances the search and browse functionalities for human users.

Ylvi provides a comprehensive user interface for capturing the mentioned *semantic glue* between media objects in the form of natural-language text articles and links between these items. Applied to the above-mentioned example, users could collaboratively create a set of articles that describe the project and its resources, and provide links to these resources. Embedded queries (queries that are included in the article source and rendered as a list of their result set) help to maintain topicality and consistency without additional effort. Such project documentation would be understandable by both computers and humans. As the underlying data model can be extended at any time, such a system could adapt to new requirements in a flexible fashion.

In this sense we consider Ylvi as a high-level, collaborative user interface for an underlying media management framework (in this case METIS) that combines the strength of a configurable semantic multimedia for data representation with the intuitive input paradigm of the emerging semantic wiki technologies. Ylvi extends other semantic wiki approaches by three main aspects: its high configurability, strong multimedia support, and adaptive semantic search.

Configurability. The underlying open architecture upon which Ylvi is based provides a broad range of configurable features:

- Configurable markup language: All syntactic elements of the markup language (except for a minimum set of core elements) can be dynamically defined/extended.
- Configurable visualisation: Ylvi articles are rendered by pipelines of rendering plugins, using the METIS cross-channel publishing framework, which supports arbitrary output channels, including XHTML, SMIL and Microsoft Word.
- Configurable semantics: All semantic modelling elements (article types, attributes, and link types) are configurable on the fly and can be introduced into the system using the functionality provided by METIS (through a Protégé interface, XML import/export, semantic packs, or the web-based GUI).
- Functional extensions: The METIS plugin frameworks for functional extensions can be used by Ylvi as well. Additional specific plugin types (e.g. render plugins, toolbox plugins) that implement Ylvi specific functionality (e.g. article rendering) were also developed.

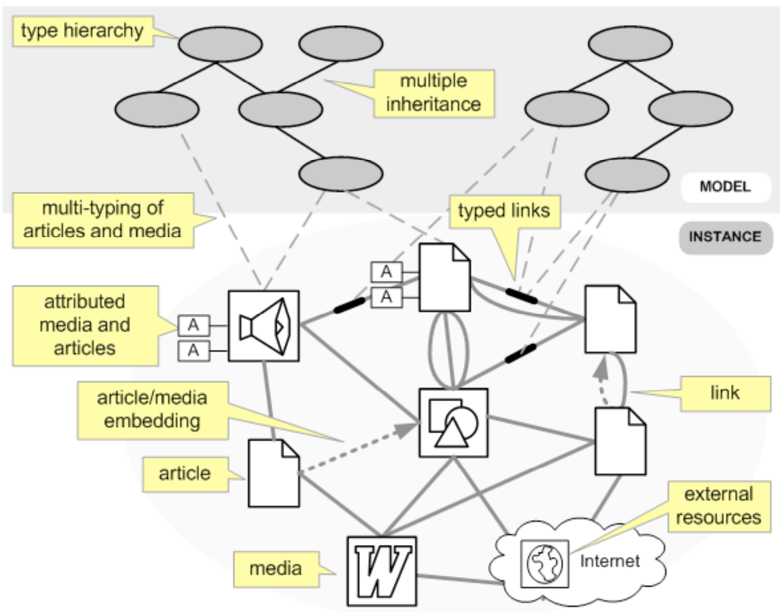

Figure 3: Semantic Features in Ylvi

Multimedia support. Ylvi can be characterised as a multimedia-specific wiki as it provides an abstract data model for both simple and complex media, transparent access to arbitrary data sources through a plug-in locator mechanism, as well as support for multimedia aggregation and multi-channel publishing. Ylvi treats both wiki articles and multimedia objects in a uniform way: Both are modelled as METIS media objects that can be typed and attributed and may participate in directed, typed links. An overview of the semantic features provided by Ylvi is depicted in Figure 3 and described in more detail below.

Semantic Search. Ylvi provides comprehensive search facilities, based on a full-text index of articles on the one hand, and on its extensive semantic features on the other hand. Articles and media instances can be related and annotated using *multi-typing, attribution*, and *typed links* (for internal and external resources). The metadata expressed by these modelling primitives can be searched by an *adaptive search algorithm* that constantly restricts search space and assists the user in narrowing down his desired search results.

> **Example** Our experience from recent research and development projects indicates that multimedia-enabled semantic wikis such as Ylvi are highly applicable for collaborative document and media management in inter-organizational projects. In the course of various research projects, we employed Ylvi as shared project management and documentation platform, whereas the project ontology was developed on-the-fly during project work, using the Ylvi interface. We consider our experience of adapting Ylvi to concrete project requirements within a very short development time as a proof of concept for (1) the rapid prototyping goals we aimed at with the development of the METIS platform, and (2) the goal of developing a Semantic Wiki framework that is easy to adapt to concrete application settings.

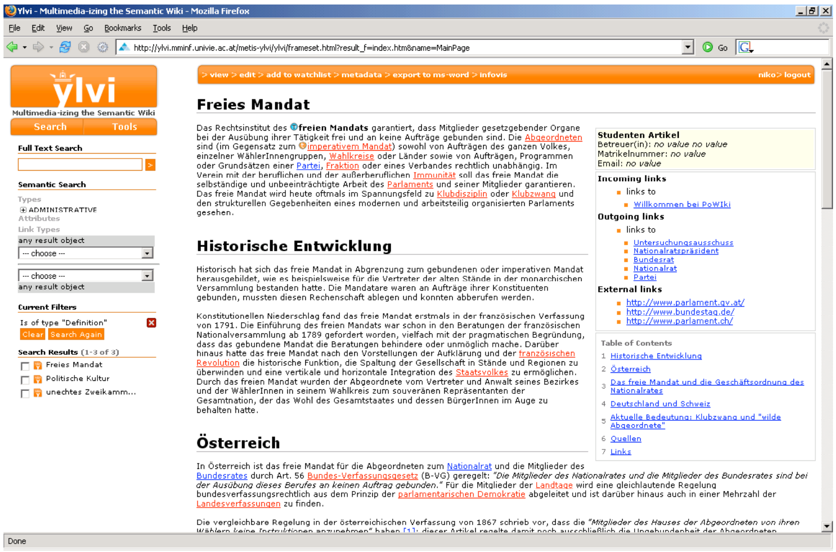

Figure 4: Ylvi Screenshot

SemDAV: Leveraging Business Knowledge to the Desktop

The SemDAV project (Schandl, 2006; Schandl & King, 2006) aims to provide a suite of technologies which enable users to work with data in whatever form in a unified way. SemDAV follows the following general design considerations that the authors have observed in everyday work.

- *Unify the user's view on data.* Users are confronted with a variety of data structures, user interfaces, classification and annotation schemes, interaction methods, and workflows. This menagerie is imposed by applications, each of which operates in its own "data world" that is in many cases not connected to the world outside the application.

  The file system, as the lowest common denominator of all data processing systems, could serve as a central point of cross-application and cross-domain data management. However, current file systems do not provide the optimal means for efficient organization of data. Instead, they are restricted to fixed hierarchies and minimal sets of metadata, most of which do not help the user to remove outdated or unnecessary data efficiently.

- *Provide well-known user interaction metaphors.* Because of the limitations of file systems, developers are forced to implement application-specific data and metadata formats, which, in turn, requires the design of application-specific user interaction and visualization metaphors. This not only forces the user to adapt to a plurality of interfaces, but also restricts semantic annotation and relation of data to features implemented by the respective applications.
- *Do not overload users with machine semantics.* To a great extent, semantics are defined by technical aspects of a system's implementation. It need not be the case that the contacts in a mail application differ from those in instant messaging software. It need not be the case that a picture sent via mail can have a textual description attached (the mail message itself), while a picture received on an USB stick cannot.

  Machine semantics are often expressed using complicated and expressive schemas or ontologies that have been developed to serve the application's need. However, many users are not willing or are not able to cope with this massive complexity. As popular collaborative services like Flickr[2] or del.icio.us[3] demonstrate, semantics can be expressed in a much easier and more user-centric manner, as long as collaborators share a certain level of knowledge – a requirement that holds true for a significant fraction of daily work.
- *Use open standards and technologies.* While the previous three design principles address mainly user-related aspects, this rule focuses on technical implementation. As history shows, data is often hidden in application-specific formats and schemas. In order to widely leverage data exchange between applications and systems, it is particularly important that systems provide representations of data in well-defined, open formats. In our view, the enormous success of XML[4] is a proof of this claim's importance, and the transition to semantically rich formats like RDF is the next, logical step to answer the open question of semantic interoperability.

---

[2] Flickr: `http://www.flickr.com/photos/tags/`

[3] del.icio.us: `http://del.icio.us/tag/`

[4] XML: `http://www.w3.org/XML/`

Example Using a collaborative system like a semantic wiki (e.g. Ylvi), co-workers can share project-related information and collaboratively improve project documentation. Web browsers provide a suitable technology for unified interfaces, independent of the platform or context from which they are used. However, there will always be the need for private data and personal information that is stored only on personal devices but can nevertheless be related to shared information, e.g. personal meeting minutes, or mail messages. The SemDAV system will be able to integrate and relate data from various sources and can help to bridge the gap between the public and private spheres.

The SemDAV system architecture for data integration is based on the client/server paradigm, and one of its core components, the SemDAV protocol, can be considered as semantic extensions for the WebDAV protocol (Goland et al., 1999). This architecture can be applied to various application scenarios, since both client and server components can be executed on the same physical machine; and a SemDAV server may act as client for another server.

SemDAV proposes a basic abstract atomic data item called a *sile* (from *semantic file*). A sile can be any digital object, a file, an image, a piece of music, a person, a machine, an e-mail message. A sile can be compared to a resource in RDF, to an object in object oriented programming languages, to a file in a file system, and to any object in the physical world. Siles are self-contained data units that can be subject to attribution, semantic annotation (tagging or classification), and association with other siles.

The main components that operate on siles are depicted in Figure 5. The *Semplorer user interface* (1) (Schandl, Amiri, Pomajbik, & Todorov 2007) is the central tool for data and metadata management. It provides a user interface that is oriented towards well-known file management utilities, like Windows Explorer. Using the Semplorer, users are able to browse and search SemDAV repositories, and to manipulate associated metadata in a simple way.

Metadata processing and manipulating is not only done by the user by the means of the Semplorer, but also through applications that are aware of

SemDAV features (2). Applications are the interface through which the user actually works with data; thus, it is particularly important that applications track and store metadata as early as possible, especially if no additional user input is required (cf. Schandl & King, 2006).

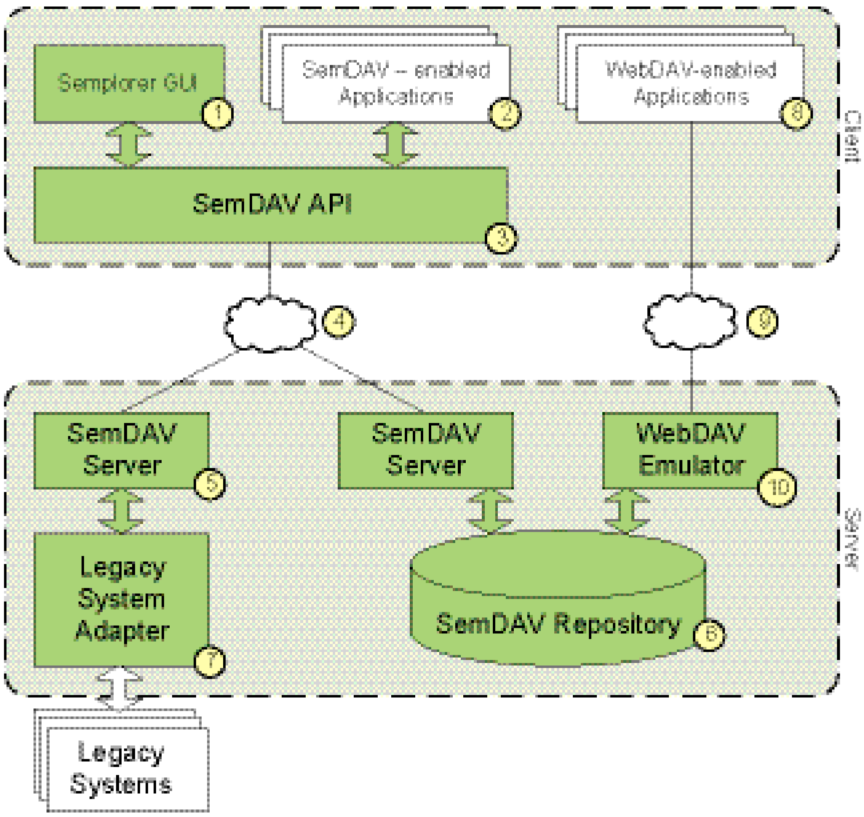

Figure 5: SemDAV Architecture

Applications, as well as the Semplorer, access SemDAV repositories through the SemDAV API (3). This API abstracts from any implementation details and provides convenient methods to access siles and for search, retrieval and modification of their associated metadata. Currently, an API for Java 1.5 is under development.

If the client and the server components are executed on the same machine, requests issued through the SemDAV API are directly handed over to the server request handler. However, if this is not the case, requests are translated into SemDAV protocol requests (4). This protocol is built on top of the standardized technologies HTTP (Fielding et al., 1999) for content handling and transfer, and SPARQL (Kendall, 2006) for metadata querying and processing. Thus, all features that can be applied to these technologies (like network transport, routing, and encryption, as well as query rewriting and optimization) can be applied to SemDAV without further effort. A SemDAV server implementation (5) handles the requests and processes them either using a designated local repository (6) or by accessing external data sources using specially developed adapters (7). A designated repository combines storage capabilities for binary data (e.g. in the file system) and metadata (e.g. using a RDF triple store). Legacy system adapters, on the other hand, may be implemented in order to access e.g. mail servers, CRM data bases, or web resources.

To gain backwards compatibility with existing systems and applications, SemDAV servers will be able to handle WebDAV requests and map them to the SemDAV data model. WebDAV is an extension to HTTP that allows clients to execute authoring tasks on remote servers. WebDAV features can be mapped to common file system metaphors (see e.g. davfs2[5]) and is supported by all major desktop operating systems, including Microsoft Windows, Apple Mac OS X, and a magnitude of Linux derivates. The WebDAV protocol is also implemented by a (perhaps surprising) number of commercial applications. This will allow applications (8) to transparently access data stored in the SemDAV server without any further modification. WebDAV requests (9) are interpreted by a WebDAV server emulation (10) and are mapped to operations on SemDAV repositories.

5. *Future Research Directions*

We identify three main trends that will significantly change the way knowledge workers interoperate in project-driven environments. First, the proliferation of semantics based on Semantic Web technologies will allow the development of more interoperable systems; an evolutionary process in which XML was only the first step. Second, collaborative knowledge organization metaphors (like tagging or classification) will gain more importance, and people will be willing to participate in such

[5] WebDAV Linux File System (davfs2): `http://dav.sourceforge.net`

systems if they can discover added value for themselves. Third, despite the blurring borderlines between local and online applications, there will always be the need for personal data; thus, applications increasingly will have to provide methods for merging remote and local resources.

Using the frameworks presented in this chapter as basis technology, we gave examples for an application suite that addresses these future challenges. METIS is able to semantically integrate multimedia information from various distributed sources; Ylvi provides a collaborative, web-based interface for rapid knowledge exchange and management, and SemDAV and its client implementation is able to integrate personal and shared semantic information and provides interfaces to quickly manage, annotate, and retrieve information.

For the future knowledge worker, we envision a seamless interface that provides a single point of entry to all the user's information needs. Old, unorganised data swamps, like file systems, mail servers, and web pages, will be wrapped and integrated into fully interconnected, search- and browsable knowledge meshes. With nearly infinite storage capacity, users will not have to worry about archiving or deleting information – knowledge organization will be performed with very little additional effort *at worktime*.

6. Acknowledgements

This work has been partially supported by the Austrian Federal Ministry of Transport, Innovation, and Technology, and the Austrian Federal Ministry of Economics and Labour. The authors also thank Bernhard Haslhofer for valuable contributions to this chapter.

7. References

Aberer, K., Cudre-Mauroux, P., Hauswirth, M., & van Pelt, T. (2004). GridVine: Building Internet-Scale Semantic Overlay Networks. In *Proceedings of the 3rd International Semantic Web Conference (ISWC)* (pp. 107-121). Hiroshima, Japan: Springer.

Anjomshoaa, A., Manh Nguyen, T., Shayeganfar, F., & Tjoa, A. (2006). Web Service Based Business Processes Automation Using Semantic Personal Information Management Systems – The Semantic Life Case. In *Proceedings of the 6th International Conference on Practical Aspects of Knowledge Management* (pp. 1-12). Vienna, Austria: Springer.

Ajchrzak, A., Wagner, C., & Yates, D. (2006). Corporate Wiki Users: Results of a Survey. In *Proceedings of the 2006 International Symposium on Wikis* (pp. 99-104). Odense, Denmark: ACM Press.

Bizer, C., & Seaborne, A. (2004). *D2RQ - Treating non-RDF-databases as virtual RDF graphs*. In *Proceedings of the 3rd International Semantic Web Conference*. Hiroshima, Japan: Springer.

Decker, S., Erdmann, M., Fensel, D., & Studer, R. (1998). Ontobroker: Ontology Based Access to Distributed and Semi-Structured Information. In *DS-8: Proceedings of the IFIP TC2/WG2.6 8th Working Conference on Database Semantics - Semantic Issues in Multimedia Systems* (pp. 351-369). Deventer, The Netherlands: Kluwer.

Fielding, R., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., Berners-Lee, T. (1999). *Hypertext Transfer Protocol – HTTP/1.1 (RFC 2616)*. The Internet Society.

Fitzgerald, M. (2003). Trash Your Desktop. *MIT Technology Review*, November 2003, pp. 42-46.

Géczy, P., Izumi, N., Akaho, S., & Hasida, K. (2006). Extraction and Analysis of Knowledge Worker Activities on Intranet. In *Proceedings of the 6th International Conference on Practical Aspects of Knowledge Management* (pp. 73-85). Vienna, Austria: Springer.

Goland, Y., Whitehead, E., Faizi, A., Carter, S.R., Jensen, D. (1999). *HTTP Extensions for Distributed Authoring – WebDAV (RFC 2518)*. The Internet Society.

Haslhofer, B. (2006). A Service Oriented Approach for Integrating Metadata from Heterogeneous Digital Libraries. In *Proceedings of the 1st International Workshop on Semantic Information Integration on*

Knowledge Discovery. Yogyakarta, Indonesia: Austrian Computer Society Book Series.

ISO (2003). *The Dublin Core Metadata Element Set*. International Standard ISO 15836-2003.

Karger, D., Bakshi, K., Huynh, D., Quan, D., & Sinha, V. (2005). Haystack: A General-Purpose Information Management Tool for End Users Based on Semistructured Data. In *Proceedings of the 2nd Biennial Conference on Innovative Data Systems Research* (pp. 13-26). Asilomar, CA, USA: Online Proceedings.

Kaulgud, V.S. & Dolas, R. (2006). DKOMP: A Peer-to-Peer Platform for Distributed Knowledge Management. In *Proceedings of the 6th International Conference on Practical Aspects of Knowledge Management* (pp. 119-130). Vienna, Austria: Springer.

Kendall, G.C. (ed.) (2006). *SPARQL Protocol for RDF (W3C Candidate Recommendation, 6 April 2006)*. W3C.

King, R., Popitsch, N., & Westermann, U. (2004). METIS: A Flexible Database Foundation for Unified Media Management. In *Proceedings of the 12th Annual ACM International Conference on Multimedia* (pp. 744-745). New York, NY, USA: ACM Press.

King, R., Popitsch, N., & Westermann, U. (2007). METIS: A Flexible Foundation for the Unified Management of Multimedia Assets. In *Multimedia Tools and Applications, 33*(3) (pp. 325-349)

Lamb, R., Davidson, E. (2005). Understanding Intranets in the Context of End-User Computing. *ACM SIGMIS Database*, *36*(1) (pp. 64-85)

Le Coche, E., Mastroianni, C., Pirrò, G., Ruffolo, M., & Talia, D. (2006). A Peer-to-Peer Virtual Office for Organizational Knowledge Management. In *Proceedings of the 6th International Conference on Practical Aspects of Knowledge Management* (pp. 166-177). Vienna, Austria: Springer.

Lethi, P., & Frankhauser, P. (2004). XML Data Integration with OWL: Experiences and Challenges. In *2004 Symposium on Applications and the Internet (SAINT 2004)* (pp. 160-170)

Maedche, A., Motik, B., Silva, N., & Volz, R. (2002). MAFRA - An Ontology Mapping Framework in the Semantic Web. In *Proceedings of the ECAI Workshop on Knowledge Transformation*. Lyon, France

Mena, E., Illarramendi, A., Kashyap, V., & Sheth, P. (2000). OBSERVER: An Approach for Query Processing in Global Information Systems Based on Interoperation Across Pre-Existing Ontologies. In *Journal of Distributed and Parallel Databases, 8(2)* (pp. 223-271)

Martínez, J.M., Koenen, R., & Pereira, F. (2002). MPEG-7: The Generic Multimedia Content Description Standard, Part 1. *IEEE MultiMedia, 9*(2), 78-87.

Nejdl, W., Wolf, B., Changtao, Q., Decker, S., Sintek, M. et. al. (2002). EDUTELLA: a P2P networking infrastructure based on RDF. In *Proceedings of the 11th International World Wide Web Conference* (pp. 604-615). Honolulu, Hawaii, USA: ACM Press

Noy, N. (2004). Semantic Integration: A Survey of Ontology-based Approaches. In *SIGMOD Rec., 33(4)* (pp. 65-70)

Noy, N.F., Sintek, M., Decker, S., Crubezy, M., Fergerson, R.W., & Musen, M.A. (2001). Creating Semantic Web Contents with Protégé-2000. *IEEE Intelligent Systems, 16*(2) (pp. 60-71)

Oren, E., Breslin, J.G., & Decker, S. (2006). How Semantics Make Better Wikis. In *Proceedings of the 15th International World Wide Web Conference* (pp. 1071-1072). Edinburgh, Scotland: ACM Press.

Popitsch, N., Schandl, B., Amiri, A., Leitich, S., & Jochum, W. (2006). Ylvi – Multimedia-izing the Semantic Wiki. In *Proceedings of the 1st Workshop on Semantic Wikis – From Wiki to Semantics*. Budva, Montenegro: CEUR-WS Vol. 206.

Priebe, T., & Pernul, G. (2003). Towards Integrative Enterprise Knowledge Portals. In *Proceedings of the 12th International Conference* *on Information and Knowledge Management* (pp. 216-223). New Orleans, LA, USA: ACM Press.

Risse, T., Knezevic, P., Meghini, C., Hecht, R., & Basile, F. (2005). The BRICKS Infrastructure – An Overview. In *Proceedings of the 75th Conference on Electronic Imaging, the Visual Arts & Beyond* (EVA 2005), Moscow, Dec. 2005

Sauermann, L., Bernardi, A., & Dengel, A. (2005). Overview and Outlook on the Semantic Desktop. In *Proceedings of the ISWC 2005 Workshop on the Semantic Desktop*. Galway, Ireland: CEUR-WS Vol. 175.

Sauermann, L. (2005). The Gnowsis Semantic Desktop for Information Integration. In *Proceedings of the 3rd Conference on Professional Knowledge Management*. Kaiserslautern, Germany: Springer.

Schandl, B. (2006). SemDAV: A File Exchange Protocol for the Semantic Desktop. In *Proceedings of the Semantic Desktop and Social Semantic Collaboration Workshop*. Athens, GA, USA: CEUR-WS Vol. 202.

Schandl, B., & King, R. (2006). The SemDAV Project: Metadata Management for Unstructured Content. In *Proceedings of the 1st International Workshop on Contextualized Attention Metdata: Collecting, Managing and Exploiting of Rich Usage Information* (pp. 27-32). Arlington, VA, USA: ACM Press.

Schandl, B., Amiri, A., Pomajbik, S., & Todorov, D. (2007). Integrating File Systems and the Semantic Web. *Demo at the 3rd European Semantic Web Conference*. Innsbruck, Austria.

Tolksdorf, R., & Simperl, E.P.B. (2006). Towards Wikis as Semantic Hypermedia. In *Proceedings of the 2006 International Symposium on Wikis* (pp. 79-88). Odense, Denmark: ACM Press.

Völkel, M., Krötzsch, M., Vrandecic, D., Haller, H., & Studer, R. (2006). Semantic Wikipedia. In *Proceedings of the 15th International Conference on World Wide Web* (pp. 585-594). Edinburgh, Scotland: ACM Press.

Völkel, M., & Oren, E. (2006). Towards a Wiki Interchange Format (WIF). In *Proceedings of of the 1st Workshop on Semantic Wikis – From Wiki to Semantics*. Budva, Montenegro: CEUR-WS Vol. 206.

Wache, H., Vögele, T., Visser, U., Stuckenschmidt, H. Schuster, G., Neumann, H., Hübner, S. (2001). Ontology-based Integration of Information. In *Proceedings of the IJCAI-01 Workshop on Ontologies and Information Sharing* (pp. 108-117).

Wen, J., Li, Q., Ma, W., & Zhang, H. (2003). A Multi-Paradigm Querying Approach for a Generic Multimedia Database Management System. *ACM SIGMOD Record, 32*(1), 26-34.