

IkeWiki: A Semantic Wiki for Collaborative Knowledge Management

Sebastian Schaffert

Salzburg Research Forschungsgesellschaft/Salzburg New Media Lab
Jakob Haringer Str. 5/II, A-5020 Salzburg, Austria
<http://www.salzburgresearch.at>

Abstract

This article describes the architecture and design principles of IkeWiki, a Semantic Wiki we developed at Salzburg Research. Outstanding features of IkeWiki are its support for collaborative knowledge engineering, its ease of use, its support for different levels of formalisation ranging from informal texts to formal ontologies, and its sophisticated, interactive user interface. While IkeWiki has been developed primarily as a tool for ontology engineering, it can be used in a variety of application scenarios. We briefly present some of these at the end of the article.

1 Introduction

A serious obstacle for the development of Semantic Web applications is the lack of formal ontologies and knowledge represented in other formal languages. Arguably, one of the main reasons for this is the rather high technical barrier for using Semantic Web technologies that deters many domain experts from formalising “their” knowledge. Therefore, formalised knowledge is currently limited to computer science and a small number of areas with a very high demand of interoperability (e.g. genetic engineering).

At the same time, wiki systems are becoming more and more popular as tools for content and knowledge management. Much knowledge is nowadays available in systems like Wikipedia. Unfortunately, this vast knowledge is not accessible for machines. If a small amount of this knowledge would be formalised, wiki systems could provide improved interfaces and advanced searching and navigation facilities.

“Semantic Wiki” systems aim to combine “traditional” wiki systems with Semantic Web technology. This combination bears much potential in many application areas. In this article, I present our prototype system *IkeWiki* (*ike* = *knowledge*, *wiki* = *fast*), which is a feature-rich semantic wiki system we developed at Salzburg Research. IkeWiki serves as a testbed for possible applications of semantic

wiki technology. I briefly outline three of these applications (*collaborative ontology engineering*, *collaborative knowledge management*, and *educational environments*).

This article is structured as follows: Section 2 summarises important features of (traditional) wiki systems; Section 3 introduces Semantic Wikis and describes features common to all or many of them; Section 4 briefly describes the design and architecture of IkeWiki; Section 5 presents the three application scenarios; Section 6 gives an overview over related work; finally, Section 7 concludes with some perspectives.

2 Traditional Wiki Systems

“Wiki” is the short form for “WikiWikiWeb” and is derived from the Hawaiian expression “wiki wiki” meaning “fast” or “quick”. A wiki is essentially a collection of Web sites connected via hyperlinks. While there is a wide range of wiki systems available (e.g. MediaWiki¹, MoinMoin², TWiki³) with different purposes and audiences, all of them share the following common properties:

Editing via Browser. Content is usually edited via a simple browser interface that can be used without installing any additional (expensive) software. This makes editing simple and allows to modify pages from everywhere in the world with only minimal technical requirements. As a consequence, content creators can access and update the wiki from wherever they are, e.g. at work, at home, at conferences, nowadays even while travelling.

Simplified Wiki Syntax. Content is usually expressed in a simplified hypertext format (“wiki syntax”) that is much easier to use for non-technical users than e.g. HTML. Formatting thus does not require knowledge of HTML.

¹<http://www.mediawiki.org>

²<http://moinmoin.wikiwikiweb.de/>

³<http://twiki.org/>

Rollback Mechanism. Changes to the content of a wiki are *versioned* each time they are stored, i.e. previous versions of pages are kept. This allows to revert back to earlier versions of a page e.g. in case important parts have been accidentally deleted or undesirable modifications have been made by someone else. Also, most wiki systems allow to compare two versions of a page, making it possible to identify changes between edits quickly.

Unrestricted Access. In most wiki systems, access is completely unrestricted – i.e. anyone can correct, modify, complete, or even delete anything. While this might seem strange, and even dangerous, from a traditional perspective, practice shows that the system works: on the one hand, ill-meaning users are rather rare; on the other hand, all changes can easily be undone using the rollback mechanism.

Note that some wikis still allow to apply further access restrictions using users and groups as found in traditional content management systems.

Collaborative Editing. The above-mentioned properties combined make wikis an ideal tool for collaborative editing. As soon as someone creates content, others can contribute to it, extend it, correct it, etc. Many wiki systems provide further support for collaborative editing, e.g. by means of discussion forums, summaries of changes, and list of last updates.

Strong Linking. Pages in a wiki are usually strongly linked with each other using hyperlinks. The reason for this is that the simplified wiki syntax makes it very easy to define a link to another page in the wiki. For example, in many wikis a link is defined by enclosing a word in square brackets, or by using a so-called “CamelCase” where a word contains several upper-case letters. Links to non-existing pages are usually rendered in a different colour. If a user clicks on such a link, the system redirects him to a view where he can create the non-existing page. In many wikis, this is even the only way to create a page.

Links in a wiki are the most important tool for navigation. Therefore, many systems allow not only to follow links in the direction they are defined but also in reverse direction (“back-links”).

Search Function. As a further tool for navigation, almost all wiki systems allow a full-text search over the content of all pages.

Uploading of Other Content. In addition to creating text content, some of the more sophisticated wikis allow uploading of arbitrary (multimedia) content, like documents, images, program code, etc.

Unlike other groupware or content/knowledge management tools, a wiki system gives users almost complete freedom over the content development process without rigid workflow, access restrictions, or predefined structures. Users need not adapt their practice to the “dictate of the system”, but can allow their own practice to define the structure. This is important, because different domains often have – or even require – different kinds of workflow.

As a recent survey on the popular technology site *Slashdot* showed⁴, wiki systems are currently used for a wide variety of purposes, including:

- *encyclopaedia systems*: collect knowledge in a certain area (e.g. Wikitravel⁵) or unrestricted (e.g. Wikipedia⁶) in a community effort with contributions from a wide range of users
- *software development*: collaboratively create documentation, collect ideas, track bugs; most of today’s high-profile Open Source projects (e.g. Apache, Mozilla, OpenOffice) use wikis for coordination
- *project knowledge management*: project tracking, brainstorming and exchange of ideas, coordination of activities, agenda tool for collecting topics of meetings, project notes repository, knowledge base, staff directory
- *personal knowledge management*: sketchpad to collect and elaborate personal ideas, addresses, dates, tasks, bookmarks, etc. [6]
- *collaborative writing*: authors work collectively on a writing (short story, novel, etc.) which is immediately accessible by readers for their enjoyment
- *CMS/knowledge base*: collect content, connect content, simple publication tool

3 Semantic Wiki Systems

A “semantic wiki” extends a wiki by “semantic technologies” like RDF, OWL, Topic Map, or Conceptual Graphs. The main idea is to make the inherent structure of a wiki – given by the strong linking between pages – accessible to machines (agents, services) beyond mere navigation. This is generally done by annotating existing navigational links with symbols that describe their meaning. For example, a link from *Mozart* to *Salzburg* could be annotated with *lived in* or *born in*.

⁴<http://ask.slashdot.org/article.pl?sid=06/01/21/1958244>

⁵http://wikitravel.org/en/Main_Page

⁶<http://www.wikipedia.org>

Such annotations are useful for many purposes, e.g. enhanced presentation by displaying contextual information, enhanced navigation by giving easy access to relevant related information, and enhanced “semantic” search that respects the context in addition to the content. Note that presentation, navigation, and search can be done in a rather generic manner, but often profit greatly from an adaptation to the represented context.

Semantic wikis exist in many different flavours (e.g. SemMediaWiki [3], SemWiki [6], IkeWiki [4], PlatypusWiki [5]). While for some the page content is still in the foreground and the annotations just optional “added value”, others require annotations – sometimes to the extent where the annotation is more important than the page content itself. This reflects also that different systems have different purposes, e.g. extending existing content by annotations to allow for better navigation, collaborative ontology engineering, etc. Nonetheless, frequently found features are:

Typing/Annotating of Links. Virtually all semantic wikis allow to annotate links by giving them certain types. The idea behind this is that a link created by a user almost always carries meaning beyond mere navigation, as given in the example in the beginning of this section. The way link annotations are edited differs from system to system. Some semantic wikis include the annotations as part of the wiki syntax (e.g. *Semantic MediaWiki* [3]), while others provide a separate editor for adding annotations (e.g. *IkeWiki*).

Context-Aware Presentation. Many semantic wikis can change the way content is presented based on semantic annotations. This can include enriching pages by displaying of semantically related pages in a separate link box, displaying of information that can be derived from the underlying knowledge base (e.g. a box with a graphical tree presentation for content belonging to a hierarchy, or license information), or even rendering the content of a page in a different manner that is more suitable for the context (e.g. multimedia content vs. text content).

Enhanced Navigation. Annotated/typed links provide more information for navigation. Whereas a traditional wiki only allows to follow a link, a semantic wiki offers additional information about the relation the link describes. Such information can be used to offer additional or more sophisticated navigation. For instance, links are more independent from the textual context they appear in and can be displayed e.g. in a separate “related information” box. The page describing Mozart could e.g. offer a separate box with references categorised by “lived in”, “composed”, etc.

Semantic Search. Most semantic wikis allow a “semantic search” on the underlying knowledge base. Usually, queries

are expressed in the language SPARQL, a query language recently proposed as W3C recommendation for RDF querying. Using “semantic search”, users can ask queries like “retrieve all pieces composed by Mozart” or “retrieve all documents where the license permits derivative works”.

Reasoning Support. Reasoning means deriving additional, implicit knowledge from the facts entered into the system using predefined or user-defined rules in the knowledge base. For example, from the fact that “Mozart” composed “Die Zauberflöte”, a system capable of reasoning could deduce that “Mozart” is a “Composer”. Although reasoning is an important feature, it is only supported by a small number of Wikis. The reasons for this might be that it is time-consuming, memory intensive, and can yield results that are not expected and/or traceable by the user.

4 IkeWiki

A number of semantic wiki systems are currently under development.⁷ In the following, we introduce one such system called *IkeWiki*. Like “wiki”, the name is derived from Hawaiian and can be translated as “fast knowledge”.

4.1 Design Principles

Although now also considered in different settings, *IkeWiki* has originally been developed as a prototype tool to support knowledge workers in collaboratively formalising knowledge [4]. *IkeWiki*’s design principles are influenced by this idea, but equally hold in other areas:

Easy to Use, Interactive Interface. As domain experts are usually non-technical people, ease of use is of big importance. Therefore, *IkeWiki*’s interface resembles as closely as possible the Wikipedia interface which people are familiar with. Also, *IkeWiki* offers an interactive WYSIWYG editor (using AJAX technology to communicate with the server backend) in addition to the traditional structured text editor, as WYSIWYG editors generally have a better acceptance among non-technical users. The WYSIWYG editor also supports interactive typing of links and resources.

At the same time, *IkeWiki* offers all the normal editing interfaces (supporting Wikipedia-style structured text) for more experienced users.

Compatibility with Wikipedia/MediaWiki. A significant amount of “informal” knowledge is available in Wikipedia. To reuse this knowledge, *IkeWiki* supports the

⁷<http://wiki.ontoworld.org/index.php/Semantic-Wiki.State.Of.The.Art>

Wikipedia syntax. This allows users to import existing content from Wikipedia into IkeWiki (e.g. via simple copy and paste) and begin with semantic annotations straight away.

Compatibility with Semantic Web standards. To be able to exchange data with other applications (e.g. ontology editors, Semantic Web services, other wikis), IkeWiki is purely based on existing Semantic Web standards like RDF and OWL. Note that other knowledge representation formats like conceptual graphs are conceivable but not investigated at this point.

Immediate Exploitation of Semantic Annotations. An important motivating aspect of wiki systems is that content is immediately available to the public when a user clicks on “save”. Similarly, IkeWiki allows immediate exploitation of semantic annotations for enhanced editing, presentation, navigation, and searching, even if the knowledge base is not yet fully formalised.

Support for Different Levels of Experience. IkeWiki is designed as a tool for collaborative knowledge engineering. In such a process, it is common that non-technical domain experts work together with experienced knowledge engineers. Therefore, IkeWiki supports all levels of experience. This means that certain advanced functionalities can be hidden from novice users but are available to experienced users.

Support for Different Levels of Formalisation. Different application areas need different levels of formalisation [4], and as Jim Hendler said:⁸ “a little semantics goes a long way”. One of the goals of IkeWiki is thus to support formalisation of knowledge all the way from informal texts to formal ontologies. Also, this means that parts of the knowledge base might be more formalised than others, and that formal knowledge is in constant evolution.

Support for Reasoning. Unlike most other semantic wikis, IkeWiki supports reasoning on the knowledge base. We consider reasoning important as it allows to derive knowledge that is not explicit in the data and is thus the true power of Semantic Web technology. At the moment, IkeWiki supports only OWL-RDFS reasoning, but an extension with a user-accessible rule engine is planned.

4.2 Architecture

IkeWiki is implemented as a Java web application using a layered architecture shown in Figure 1. Data is stored in a

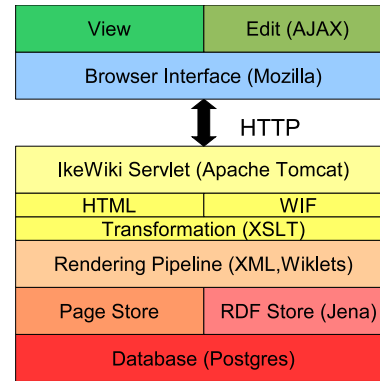


Figure 1. IkeWiki architecture (cf. Section 4.2)

Postgres database. When a resource is requested, the XML page content and related RDF data are retrieved and combined in the RenderingPipeline into an enriched XML representation. This XML representation is then either offered as interchange format for other Web services or transformed into HTML for presentation in the user’s browser. The individual layers are described in more detail in the following:

Page Store. The page store component serves to store and retrieve page content from the database. It also supports versioning and restoring of previous revisions. Page content is represented in structured text as well as an XML format we call *WIF* (wiki interchange format). Basic WIF merely describes the page content and structure, but allows to add custom, application specific information in separate namespaces. We have proposed WIF as a format for interchange of page content between arbitrary wiki systems.

RDF Store. The knowledge base is represented using the Jena RDF framework.⁹ IkeWiki uses an in-memory model to ensure responsiveness of the system. The in-memory model is frequently synchronised with a database model for persistent storage. Part of the RDF store is a SPARQL engine that allows for searching of the knowledge base.

Rendering Pipeline. The Rendering Pipeline combines page content with semantic annotations. Its output is a WIF document enriched by relevant semantic annotations (e.g. link types, context adaptation) in RXR format [2]. The “pipeline” consists of small steps we call “wiklets”. Each wiklet adds a specific piece of information to the WIF document. Wiklets can be enabled and disabled and associated with permissions so that only selected users can see the added information.

⁸as conference chair in the opening speech of the 2003 International Semantic Web Conference; Sanibel Island, Florida, USA, October 2003

⁹<http://jena.sourceforge.net>

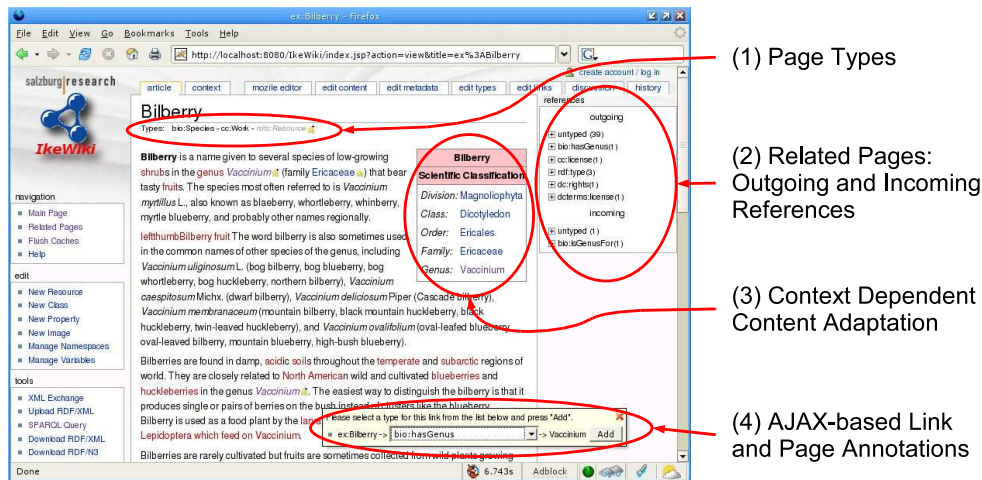


Figure 2. IkeWiki interface (cf. Section 4.3)

Transformation, Servlet. The transformation layer applies XSLT transformations to the enriched WIF format to generate XHTML, SVG, and/or WIF output. The IkeWiki servlet then serves this output to the user's browser.

4.3 Interface

IkeWiki uses a purely browser-based interface (cf. Figure 2). The current implementation only supports the Mozilla browser family due to its standards compliance and free availability.

Page View. A sample page view is shown in Figure 2. In the figure, you can see a sample article (copied from Wikipedia) about the "Bilberry". Type information is shown below the page title (1). Links to (semantically) related pages are displayed in a separate "references box" on the right hand side (2). The taxonomy box (3) showing the biological classification of the described plant is automatically generated from existing semantic annotations (i.e. Bilberry *hasGenus* Vaccinium) and is an example for context adaptation. Finally, (4) shows interactive typing of links using AJAX technology.

Content Editor. The content editor is available in two flavours: as a traditional structured text editor and as a WYSIWYG editor. The structured text editor is aimed at expert users that are familiar with other wiki systems, and allows to directly copy content from Wikipedia. The WYSIWYG editor is aimed at novice users creating new content. The WYSIWYG editor interacts with the server backend: links are automatically recognised and verified, and semantic annotations can be done directly in the editor (as also shown in Figure 2, (4)).

Semantic Annotations Editor. Semantic annotations are separated into three editors: the *metadata editor* allows to fill in textual metadata related to a page (like Dublin Core metadata or RDF comments). The *type editor* allows to associate one or more of the types available in the system with a page. The *link editor* allows to annotate outgoing and incoming links with type information. In the editors, available annotations are determined by the reasoner based on the page and link types; for example, if a link from "Mozart" to "Die Zauberflöte" is annotated by "composerOf", the system will automatically associate the type "Composer" with the page describing "Mozart", and this type cannot be deleted directly by the user.

5 Semantic Wiki Usage Scenarios

A wide range of application scenarios for semantic wikis are conceivable. In the following, we briefly present three areas that might benefit from semantic wikis, *ontology engineering*, *knowledge management*, and *educational environments*.

Ontology Engineering. In ontology engineering, domain experts and knowledge engineers work together to create a formal ontology. Semantic wikis can support this process very well:

- non-technical domain experts are offered an easy way to enter their knowledge;
- domain experts and knowledge engineers can collaboratively work on the ontology, each bringing their personal expertise; and

- the knowledge can be formalised in an evolutionary process, beginning with informal texts.

As an example, consider a historian who wants to create a semantic portal about Mozart. He begins filling the semantic wiki with texts, images, and maybe even audio content he has available via an easy to use editor. He can also describe simple relations between content. At some point, his experience in Semantic Web technologies does no longer suffice, so he brings in a knowledge engineer. The knowledge engineer – without in-depth knowledge of the topic – can take the simple relations and further formalise them e.g. to form an ontology. At the same time, the historian can continue filling the system with more content.

Knowledge Management. In our knowledge-based economy, knowledge management is becoming more and more important. Knowledge management means easy and free authoring and content creation on the one hand (as one did on sticky notes, etc), and structuring knowledge for easy retrieval and for finding related information one was not aware of on the other hand. Also, sharing of knowledge is an important issue. While traditional wikis already support much of this, semantic wikis offer advanced searching and navigation capabilities as described above. They are therefore very well suited for knowledge management. Note that [6] describes in detail the application of semantic wikis for knowledge management.

Educational Environments. In educational environments, semantic wikis can serve many different purposes. This paragraph focusses on two of them, *content creation* and *learning*.

Content Creation. In today’s knowledge society, learning content needs to quickly adapt to the specific needs of learners (“just in time learning”). Creation of appropriate content is, however, usually a time consuming task. A semantic wiki can be very helpful for content creation. The wiki represents an easily modifiable “learning content pool” where content creators can author, search for, and recombine learning material to form learning content tailored towards the learner’s needs. Semantic annotations can be used to facilitate retrieval of related content, to “match” learning material in a course, and to create the structure for a new course, which can then be exported e.g. to SCORM.

Learning. In constructivist and connectivist learning theories, users learn by constructing knowledge in their minds based on their experiences. Also, learning – especially of adults – is nowadays often at least partly self-directed. Semantic wikis are a good tool to support such learning, as they allow learners to collect, connect, and

structure information that they consider relevant, and to collaboratively collect and share information in a learning group. Also, they allow easy supervision/support by a coach.

6 Related Work

In the last year, a lot of semantic wiki systems with sometimes very different purposes and functionalities have been proposed, e.g. [6, 3, 1, 5]. The first such system was probably PlatypusWiki [5], which is focussed on the creation of RDF data. A constantly updated list of systems can be found at the Semantic Wiki Interest Group website.¹⁰ As a response to the growing interest, the Semantic Wiki workshop (SemWiki06), co-located with ESWC’06 in Budva, Montenegro, will take place for the first time.

7 Perspectives and Conclusion

In this article, we presented IkeWiki, a feature-rich semantic wiki system. Future directions for the development of IkeWiki might be to provide more support to knowledge engineers through linguistic processing techniques, identification of similar concepts through automated structure analysis, and guided questions based on underlying ontological knowledge. Also, it is intended to include a rule engine that allows users to define custom rules for the knowledge base.

On the application side, we are currently working on several scenarios for the use of semantic wikis in elearning environments, and we are investigating the use of IkeWiki as part of a methodology for collaborative knowledge engineering developed in the project Dynamont.

References

- [1] D. Aumueller and S. Auer. Towards a Semantic Wiki Experience – Desktop Integration and Interactivity in WikSAR. In *Semantic Desktop Workshop 2005 at ISWC’05*, Galway, Ireland, 2005.
- [2] D. Beckett. Modernising Semantic Web Markup. In *XML Europe 2004*, Amsterdam, The Netherlands, 2004.
- [3] M. Krötsch, D. Vrandečić, and M. Völkel. Wikipedia and the Semantic Web - The Missing Links. In *Proceedings of the WikiMania2005*, 2005.
- [4] S. Schaffert, A. Gruber, and R. Westenthaler. A Semantic Wiki for Collaborative Knowledge Formation. In *Semantics 2005*, Vienna, Austria, November 2005.
- [5] R. Tazzoli, P. Castagna, and S. E. Campanini. Towards a Semantic WikiWikiWeb. In *3rd International Semantic Web Conference (ISWC2004)*, Hiroshima, Japan, 2004.
- [6] M. Völkel and E. Oren. Personal Knowledge Management with Semantic Wikis. 2006.

¹⁰<http://wiki.ontoworld.org/index.php/Semantic-Wiki.State.Of.The.Art>