

KiWi Annual Report

KiWi Website: <http://kiwi-project.eu>

KiWi Wiki: <http://wiki.kiwi-project.eu>

KiWi Blog: <http://planet.kiwi-project.eu>



The objective of the project KiWi is to develop an advanced knowledge management system that makes use of Semantic Web and Social Software technologies to enable better knowledge sharing, collaboration, and information integration. To this aim, KiWi implements a social software platform based on a Semantic Wiki that can be easily customised to the requirements of the respective domain. Particularly, this platform integrates advanced semantic technologies in the areas of reasoning, reason maintenance, information extraction, and personalisation. The KiWi system will be evaluated in two concrete use cases at the partners Logica (project knowledge management) and Sun Microsystems (SunSpace community platform).

Summary of Activities

In 2009, KiWi has executed activities and made achievements in the following areas:

- the implementation of the KiWi core system, serving as a flexible technological basis for developing custom semantic social software applications; the KiWi core system has been released regularly as Open Source Software
- conceptual framework and first implementations of the KiWi enabling technologies (i.e. reasoning, querying, information extraction and personalisation) as part of the KiWi core system
- start of the implementation and integration of the two KiWi use case applications on top of the KiWi core system
- implementation of several additional projects (TagIT, InterEdu) on top of the KiWi platform

Throughout the next year we will release a first feature-complete version (1.0) of the KiWi system and evaluate the usefulness of our technologies in the two KiWi use cases, as well as the additional projects building on top of KiWi. The TagIT system building on top of KiWi will go into live operation starting 2010.

Work Area: KiWi Core System

In 2009, the main parts of the KiWi architecture and core system have been implemented and released as Open Source Software in regular releases starting March. The KiWi software architecture is a component-based service oriented architecture building on top of the Java Enterprise Edition and JBoss Seam. The main achievements in the core system so far are:

- a common data model for different kinds of social media applications that allows to represent versatile content inside the KiWi system; KiWi represents different kinds of content using the same model (e.g. wiki articles, news, persons, locations) and is thus very flexible and easy to adapt to different settings and applications

- a unified data access service for storing and retrieving content and meta-data from the KiWi system featuring extensive transaction and version management and advanced searching and querying functionalities for both content and meta-data
- core social media functionalities like (semantic) tagging, commenting, activity logging, etc., as well as advanced annotation functionalities like RDFa and thesaurus management
- integration with various other social media and content services for importing information, e.g. from Facebook and Twitter, from RSS feeds, from GPS tracks or from RDF data sources
- implementation of several different applications and user interfaces on top of KiWi, in particular the KiWi Wiki, the Dashboard, and the applications TagIT and InterEdu (see below)

The KiWi core system is currently under very active development with a regular release cycle. It can be downloaded from the KiWi download site.¹ A first feature-complete version will be released in mid-2010. During the next year, we will generalise the KiWi user interface components (advanced visualisations) and turn them into widgets that can be used also in external applications, e.g. to bring KiWi functionality into an existing wiki installation.

Work Area: KiWi Enabling Technologies

One of the main activities throughout the project is the furthering of applied research in the areas of reasoning, reason maintenance, information extraction, and personalisation, and the integration of these technologies in the KiWi core system. By the end of 2009, first implementations of the KiWi enabling technologies have already been integrated into the KiWi core system:

- *reasoning* using a rule-based forward chaining algorithm is currently capable of materialising inferred triples in the KiWi triple store based on administrator-defined rules; the next improvement will be a more efficient and less naïve forward chaining that also takes into account higher-level concepts than just triples
- *reasoning maintenance* information is tracked and will in the next implementation of the reasoner provide justifications about why the system exposes a certain behaviour in situations involving reasoning, particularly personalisation
- *information extraction* is currently used to support tagging, the annotation and markup of text content, as well as for statistical recommendations of similar articles inside the KiWi system; the information extraction functionality will be further improved throughout the next year and exposed in a more user-friendly manner
- *personalisation* is currently used to calculate personalised recommendations for users using a multifactor recommendation algorithm that takes into account the current context as well as the user model kept by the system; personalisation is currently available in related recommendations for each content item, in the personal dashboard, and in a personalised search

Throughout the next year, the enabling technologies will be integrated more tightly with the KiWi core system.

¹ <http://kenai.com/projects/kiwi/downloads>

Work Area: KiWi Use Cases

The two KiWi Use Cases evolved during the year 2009 and one of them (software knowledge management at Sun) had to be redefined slightly from the original plan. The two use cases now cover the following scenarios:

- *project knowledge management*; the idea of this use case is to be able to transform back and forth between unstructured textual descriptions of artefacts related to a project and structured data representations of these artefacts (e.g. a quality management form); the current implementation demonstrates how this can be achieved using RDFa and Semantic Forms in KiWi
- *intranet expert finder*; the software knowledge management use case has been redefined to integrate Sun's community equity algorithm into KiWi and use it within the SunSpace intranet as a means to identify which information is relevant and who is expert for which topic; in 2009, the integration of community equity in KiWi is mostly finished and we are working on a common KiWi/CEQ deployment in intranets

Both use cases will start their evaluation phase at the end of 2009. Evaluation results are expected mid of 2010.

Work Area: TagIT and InterEdu

In addition to the use cases developed as part of the KiWi project itself, we have built several more applications on top of the KiWi core system:

- *TagIT* is a social media system that integrates news articles from the local newspaper, user tips, photos, blog posts, persons, and Twitter and Facebook status updates and geolocates and displays them on a map; the idea behind this application is that users are typically interested in "stuff around them", be it professional news, user-generated content, or other users. The TagIT project is developed in collaboration with Salzburger Nachrichten and will go into live operation starting 2010.²
- *InterEdu* is a prototypical platform for educational media that allows teachers to collaboratively organise and share teaching material; InterEdu will be evaluated end of 2009 and serve as a proof-of-concept for the project partner, Education Highway.

In addition to these two applications, we are now increasingly building further projects on top of the KiWi core system, indicating the usefulness of the provided services for social media systems.

User Involvement, Promotion, and Awareness

Despite being one of the smaller projects, KiWi already has a very high visibility in the Semantic Web community. This is evident e.g. by the Best Demo Award the KiWi demonstration was able to achieve at this year's European Semantic Web Conference. In addition, the KiWi community organised and participated in a number of events for both the scientific and the professional community:

² <http://tagit2.salzburgresearch.at>

- the 4th Semantic Wiki Workshop at the European Semantic Web Conference in Heraklion, Greece, June 2009; several KiWi papers have also been presented at the workshop
- the I-Semantics conference co-located with the I-Know (knowledge management) conference in Graz, Austria, September 2009
- the 1st KiWi programming camp took place in March 2009 in Salzburg with developers from both inside the KiWi project and from outside.

Beyond the organisation of workshops, KiWi members have presented numerous peer-reviewed publications at scientific conferences and workshops in the areas relevant to KiWi.

In addition to scientific and professional events, the TagIT system building on top of KiWi has also been presented at Austria's "Researchers Night" ("Lange Nacht der Forschung") in November 2009, where it was very well received by the interested public.

In the area of software development, a small Open Source Community has already started to form with several external persons already contributing to the KiWi core system.

Future Work or Exploitations Prospects

In the coming year, the three major strands of activity in the KiWi project are the further development of the KiWi core system into a platform that can serve as the foundation for many different kinds of semantic social software systems, the finalisation of the integration of the advanced semantic technologies in the areas of reasoning, reason maintenance, information extraction, and personalisation, and evaluation in the two use cases.

It is our aim to release a first fully-functional version of the KiWi system in mid-2010 as Open Source software for interested developers to use.

