OURWEB – TRANSPARENT GROUPWARE FOR ONLINE COMMUNITIES

Miikka Miettinen[1], Jaakko Kurhila[2], Petri Nokelainen[1], Henry Tirri[3]

[1]*Helsinki Institute for Information Technology*
PO Box 2700, 02015 HUT, Finland
mmmietti@cs.helsinki.fi

[2]*Dept of Computer Science*
PO Box 84, FIN-00014 Univ of Helsinki, Finland
kurhila@cs.helsinki.fi

[3]*Nokia Research Center*
P.O. Box 407, FIN-00045 Nokia Group, Finland
henry.tirri@nokia.com

ABSTRACT

OurWeb is a system designed for on-line communities to foster engagement into purposeful activities associated with collaboration. In this paper we present the features and functionality of OurWeb, and report the experiences acquired from a university course, where the course structures were designed to emulate an on-line community and the students were required to act as members of a learning community. The experiences from the course suggest that enhancing interactions between the community members using novel tools is fruitful.

KEYWORDS

Online community tools, joint annotation, Wiki, project collaboration

1. INTRODUCTION

According to Bereiter (2002), the primary concern of a learning community is the development of *conceptual artifacts*. Bereiter's theory emphasizes the character of ideas, methods and perspectives as tools, which together constitute the community's established way of thinking about the domain. By creating, refining and communicating relevant knowledge, the members of the community participate in deliberate efforts to extend and improve the shared toolbox of conceptual artifacts. This requires continuous learning, because the results of the work immediately affect the standard being strived for. In order to contribute, individuals and work groups need to stay at the cutting edge of the community's constantly expanding knowledge.

Thinking of knowledge as an object that can be developed and shared systematically provides a useful perspective to the design of groupware applications. In particular, recognizing that the users are involved in specific goal oriented activities provides a basis for going beyond generic communication tools towards more focused technology. This kind of a step is necessary in order to address the focused needs of learning communities adequately.

The activities of a learning community are *progressive* and *open-ended* by nature (Bereiter 2002; Wenger 1998). They are driven primarily by questions and problems identified by the community members. The specific direction and scope of the activity is not known in advance, but emerges from the individual efforts and interactions of the participants. Communities typically consist of smaller groups of frequently collaborating individuals. In order to provide opportunities for fruitful cooperation between groups, it is essential to support overall awareness of the community's activities and enable multiple forms and levels of participation.

On the basis of these and a few other observations, we have identified *openness* and *transparency* as two essential design principles for Web-based groupware. By openness we mean that restrictions with regard to information sources and the applicability of the tools should be avoided as much as possible. Being engaged in question-driven and progressive efforts, learning communities should be enabled to work effectively with any available material rather than being restricted within the bounds of fixed document collections. Transparency implies that the community members should be able to *see* the activities taking place in the system, *benefit* from the activities, and *participate* in meaningful ways.

As an example of a system demonstrating these principles, the next section introduces OurWeb. It is based on our earlier tool named EDUCOSM (Miettinen et al. 2003), and contains many of the same features along with a number of additional ones. Both EDUCOSM and OurWeb have been tested extensively in the context of Web-based courses. In section 3 we describe an experiment, in which a group of Computer Science students was arranged to work like an expert community. Although this does not eliminate the need for evaluations with real online communities, some useful experiences and insights were gained. The results of the experiment are presented in section 3.2, and section 4 concludes the paper.

2. DESCRIPTION OF OURWEB

2.1 Main features

OurWeb consists of an integrated set of tools for collaborative knowledge building. The system is "invisible" to the user in the sense that it appears only as a custom popup menu that is available on any page accessed through the system. The menu is used for both handling the document that is currently visible and for navigating between the various parts of the system. In this section we describe the use of OurWeb from the community member's point of view.

2.1.1 Document pool

Learning communities need spaces for processing information that is relevant to their collaborative efforts. In practice, one of the main activities is searching for and sharing background material, and making own work available to others. While logged into OurWeb, a community member can link any potentially useful Web pages to the *shared document pool* of OurWeb. The procedure is deliberately made simple by enabling the user to just evoke the popup menu and select "Add to document pool" while the document is visible in the browser window. Adding a document to the document pool means that it becomes available to the entire community through the various index pages and the internal search engine of the system.

2.1.2 Joint annotation and discussions

The primary means of collaborating on the contents of the document pool are *joint annotations* and *threaded discussions*. These allow the members of the community to engage in *artifact-centered discourse* (Suthers 2002), in which the contributions of the community members appear in the immediate proximity of the relevant information.

Two different types of joint annotations are supported: highlights and comments. Highlights can be applied to marking important parts of the text, analogously to the way people highlight lines of text on paper. In practice, the user selects the desired text fragment with the mouse, brings up the popup menu and chooses the "Highlight" option from the menu. When attaching a comment to a fragment of text, the procedure is the same except that the user writes the desired comment in a popup window. A comment appears as a tooltip when the mouse pointer is placed on top of the commented text fragment (Fig.1).

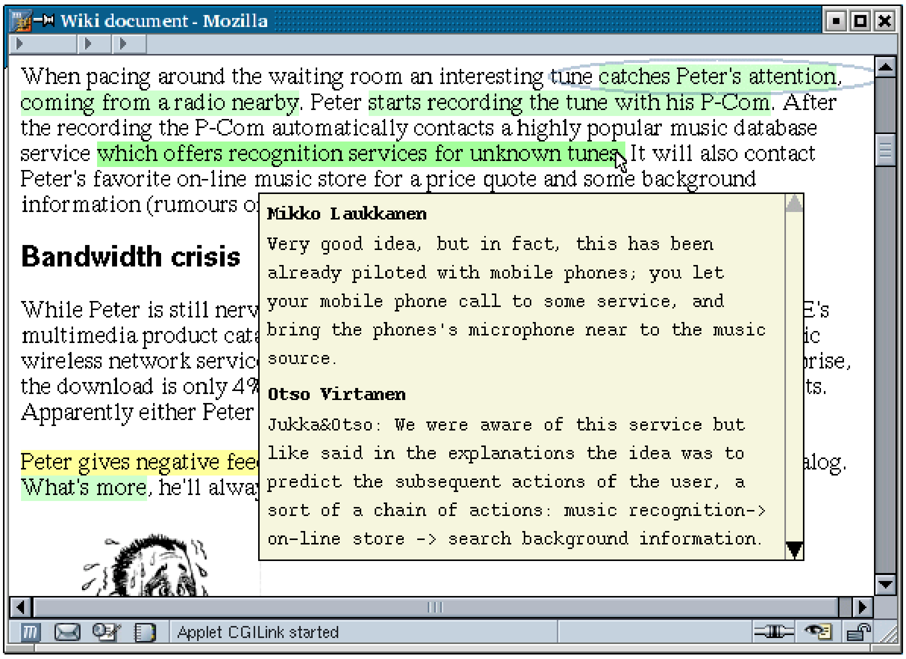

Figure 1. OurWeb comments in a tooltip

When someone else writes a comment to the same fragment (or part of the same fragment), the comment is bundled to the same tooltip when viewing the comment. This enables discussions coupled tightly to the context of the underlying text. This approach has its advantages, since it allows an easy and straightforward way to interact on details in the text (Nokelainen et al 2003). In addition, the layout or other information of the underlying document is not altered, unlike in many other annotation systems, such as Web Annotator (Reed & John 2003). Yet the annotations in OurWeb are visually easy to recognize.

Should there be longer outbursts of one's opinions, threaded discussions are available in document-specific newsgroup-type forums, which any user can attach to any document as needed. However, according to our earlier studies (Kurhila et al. 2003a; Kurhila et al. 2003b), users prefer highlights and especially tooltip comments over threaded discussion forums.

The ability to make annotations when reading has proven to be an effective strategy when trying to learn from text. According to Lonka et al. (1994), systematic highlighting was the best way to support the forming of the "big picture" from the material read. When evaluating and applying the information presented in the text, writing the text contents in one's own words was one of the most efficient strategies. Both of these annotation features --- typically used with pen and paper --- are included in OurWeb. Moreover, OurWeb brings these features to a collaborative web-based environment and extends the usefulness of these strategies by exposing the results for the benefit of the others as well.

Although OurWeb is designed to be an open environment so that users can benefit from each other's work, it is necessary that (at least some of) the annotations can be switched off as needed. This need can be addressed with *filters* that define the group of people and the time period from which annotations are shown. For example, a user may choose to view only the latest annotations of his or her own work group, or hide all the annotations completely. Filters are created in a popup window, where a community member can simply select the people from a list and optionally limit the time period to the last day, week or month.

2.1.3 Project pool

To enhance the ways to collaborate for an on-line community, OurWeb includes the possibility to suggest and carry out *projects* within the system. From the popup menu, the user is taken to a project work view, where new projects can be suggested. Every member of the community can suggest projects by filling in the fields for project title, short description and some plans or schedules. After that, anyone can join in the project. It should be noted that the project initiator does not need to join the project he or she suggested.

When someone else views the project pool, every suggested project shows title, description, proposed by -information (name and time stamp), project status (open or closed), project team members, and links to work published during the project. More information -link is available, which reveals plans and schedules and links to the background material of the project, if available.

OurWeb's project pool allows people to act in various roles. Project initiator or innovator can be a project polisher in another project. The approach taken is in harmony with the concept of communities of practice (Wenger 1998). For example, some people can be *knowledge brokers* in OurWeb and transmit information between projects exploiting the open project pool, exposure to published drafts and background materials, and easy-to-use and straightforward interaction with other projects by joint annotation.

2.1.4 Collaborative writing in Wiki

When a project is started and a team has formed, OurWeb project work view includes a Wiki to author joint projects. Wiki (or WikiWikiWeb) means that web sites and pages can be authored collaboratively with a simpler-than-HTML markup language using just a web browser as an editor (Leuf 2001). As an anecdote, geographically dispersed production of this paper has been carried out as a project in OurWeb and authored in OurWeb's Wiki. The Wiki in OurWeb is a modified version from an open-source Wiki called MoinMoin (http://moin.sourceforge.net). The authoring of documents using the modifications includes the ability to include formulas in LaTeX format and to upload images.

When the project team works with Wiki, they have an internal working copy of the project not visible to people outside the team. When appropriate, the team can publish drafts to the whole community. Either way, joint annotation features can be used upon the documents produced with the Wiki. Typically, the feedback from other community members guides the authoring process and steers the projects into more meaningful directions.

Because of the Wiki, joint annotation and overall openness and transparency, OurWeb supports the concept of *process writing* (Grabe & Kaplan 1996; Tompkins 1990) as a purposeful activity. According to Gardner and Johnson (1997) there are different stages in process writing including generating ideas for writing, rough drafting, sharing with peer reviewers, revising, editing and publishing.

2.1.5 Footprint information

Every link to a document in the document pool or published document in the project pool is attached with footprint information (see Fig. 2). The link anchor is followed by an icon of footprints. Footprint information is personal in a sense that the icon is black if there are some new highlights, comments or postings attached to the document behind the link for a particular user. If not, the icon is dimmed.

When a user places the mouse pointer over the icon, a tooltip appears with variable coloured bars for relative reading time for the document, highlights made, comments written and discussion postings made. Unseen highlights, comments and postings are visible on the bars. If the user places the mouse pointer over a bar in the tooltip, another tooltip appears stating the exact numbers (e.g. "4 new comments"). Reading time is presented as a number between 0 and 10.

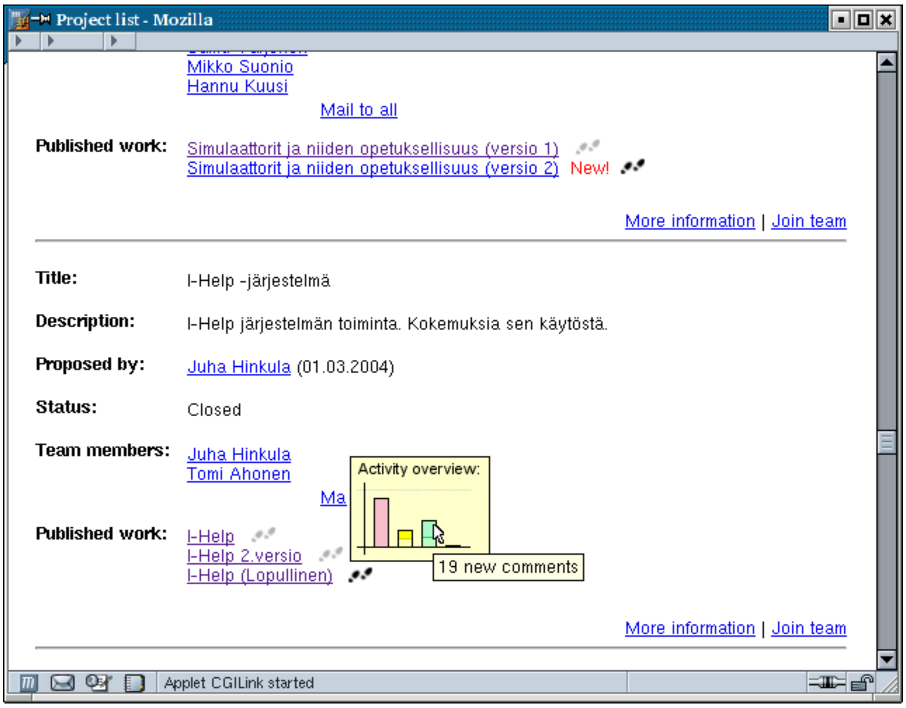

Figure 2. Footprint information of a project's published document

2.1.6 Internal search and convenience features

Search allows the user to perform a keyword search on annotations, discussion postings, or documents. In addition, the user can search the entire Web by using Google search engine.

Other features in OurWeb include a personalized desktop, multiple views to document pool, marking of new comments to help in locating annotated text with unseen comments (circled annotation in Fig.2), and email.

The desktop serves as the entry point to the system, and contains links to announcements, assignments, articles and threaded discussion messages. Personalization (recommendations) on the desktop is explained in (Miettinen et al. 2003).

2.2 Architecture and implementation

The architecture of OurWeb has gone through a number of revisions since the introduction of the system under the name EDUCOSM (see Miettinen et al. 2003 for details on the first implementation). In particular, much functionality has been moved from the browser to the server in an attempt to improve reliability and reduce the amount of browser-specific JavaScript code. New features have also been added, and the focus has moved from online courses to supporting the less structured activities of online communities.

The server acts as a proxy between the user's browser and the Web, allowing any potentially useful resources to be added to the shared document pool and augmented with annotations and other information relevant to the community or the individual user. The documents are not stored permanently on the OurWeb server. The current implementation gets some performance benefits from temporary caching, but the basic idea is that when a user requests a page, the server reads it from its original location and inserts the annotations, the popup menu, and other application specific data before forwarding it to the client. Thinking of the system as an application proxy is natural, since our goal is to provide unlimited access to any relevant parts of the Web.

All other data besides documents coming from other Web sites is stored in the file system of the OurWeb server. The application itself consists of a set of request handlers written in Python and attached to the Apache Web server with the mod_python module. We use a threading version of Apache, and in order to allow the threads to rely on shared memory, we have configured it to run in a single process. Especially document and annotation data need to be accessed very often to compute the various index pages and the footprint visualizations. Using threads with shared memory is a simple and convenient way to reduce disk usage and repetitive computations, but an embedded database would probably provide better scalability. There is no obvious reason, however, why even the current implementation would not scale easily to support communities with hundreds of active members.

Although the OurWeb server is conceptually an application proxy, its role as an intermediary is based on manipulation of URL references rather than customized browser settings. When a certain page passes through the server, one stage in the processing is to change the URLs so that they do not point directly to their real targets, but use a particular routine on the OurWeb server to retrieve the content. The routine can then supplement any linked pages with the application's additional data and functionality, and repeat the URL manipulation procedure to ensure that further requests are also directed to the proxy. The reason why we have ended up with this kind of self-made solution is that users who are behind a firewall are typically required to use a proxy to get outside the local network, and settings available in normal browsers do not allow several proxies to be chained together. Dynamic content and the ill-defined nature of HTML used in real-world pages have caused occasional problems, but the majority of them appear to be solvable with sufficiently elaborate algorithms.

The display of the content and interaction with the user are handled by a standard Web browser. All of the custom features are based on HTML and JavaScript. Annotations and the popup menu are HTML elements, whose visibility, appearance and location is controlled by JavaScript event handlers. The social navigation bar charts are dynamically generated images with associated image maps to provide the bar labels. Most of the basic functionality can be implemented using standard techniques that work reasonably well in almost any modern Web browser. The creation of new annotations is a notable exception, however, as it requires working with text fragments selected with the mouse. This is not possible in all browsers, and in any case requires some browser specific code. JavaScript based updating of document fragments was also a persistent source of occasional problems in earlier versions of the system, but nowadays we simply refresh the entire page from the server cache whenever a new annotation is made. Despite having moved as much of the work to the server-side as possible, several practical problems still prevent OurWeb from being usable with any (reasonable) browser. Fortunately, it appears that the users value OurWeb's functionality and ease-of-use enough to be willing to get themselves a suitable browser, and restricting "official" support to newer versions of Microsoft Internet Explorer, Netscape Navigator, and Mozilla has not been a major problem in practice.

An important aspect of OurWeb from the research point of view is its ability to gather significant amounts of detailed data on people's activities. Our general principle has been to log almost everything that is observable, while of course respecting the users' privacy in handling the data. In addition to the standard events related to navigation and active manipulation of contents, our logs contain e.g. information about switches of focus from the browser window to another program and back, scrolling of the page, and the display of annotations. In fact, more difficult than detecting these things (with JavaScript event handlers) is sending data almost continuously to the server. We currently use a Java applet for this purpose. It either creates a permanent connection or, if the user is behind a tight firewall, sends the data in larger chunks with the http post method. The applet is used only for logging. Critical data (e.g. annotations made by the user) is sent with an invisible form added on each page. Despite this, it seems appropriate to reduce the dependence on Java even further with browser-specific scripting, as Java has been changing from a standard feature to an optional plug-in during the last few years.

3. EMPIRICAL EVALUATION

3.1 Study setting

In practice, it is a long way from a set of tools in a system to actual well-functioning on-line communities and communities of practice. Given the natural restrictions, meaningful ways of evaluating the usefulness of the system are limited. To gather some evidence, OurWeb was used as an environment for a university-level

Web-course in Spring 2004. However, the course structure was purposefully designed to emulate the structures of an on-line expert community.

Learning communities as a paradigm for university courses takes the concept of *student-centred learning* (Felder & Brent 1996; 2001) to a new level. The responsibility for *learning* is heavily shifted towards the learner by creating greater opportunities for learning. At the same time, the learner should take responsibility for his or her own learning, but accept the responsibility for the community as well. Support for learning is not provided merely by the instructor but the community as a whole.

The course was structured for utmost freedom and deliberately fought against the common knowledge for organizing successful Web-courses. There were no face-to-face meetings apart from the starting get-together, during which the system was briefly demonstrated and the grading and other administrative issues explained. No lectures were held, no pre-made learning material was assigned or references pointed out, and no timetables were introduced except the starting date and the ending date. In particular, there was no fixed assignment cycle. The students had to come up with their own projects and take full responsibility for the initiating, organizing and scheduling the work during the course. It should also be noted that the role of the instructor in this kind of a setting is not that different from the role of a student: both influence the activity by participating in the activities within the community. The students were informed that the comments or other activity by the instructor is not to be taken as facts to be memorized; instead, the instructor should be questioned as any member of the community.

Although student-centred learning leaves a lot of the responsibility to the students, it is done primarily to achieve deeper learning and better, more applicable learning outcomes (Felder & Brent 2001). Allowing the students to choose the topics of their projects e.g. amounts to *negotiation of the domain* (Wenger 1998), provides the basis for identity transformation with its motivating side effects, increases commitment, forces the students to think about the "bigger picture", and exercises skills needed for effective functioning in the information society.

The title for the course was "Computer Uses in Education". Only 19 students were active from the start, out of the 28 who created an account for the system. The students were mostly Computer Science majors, 5 female and 14 male. The course lasted for 10 weeks, and the students were expected to invest 160 hours of work for the whole course. There were no exams. To pass the course, the students were obliged to take part in projects for ca. ten weeks and be active in the community. There were no definitive instructions for conducting the projects. Instead, it was made explicit that the projects should together paint the picture of all the important aspects of "computer uses in education". It was said that in a way — at the end of the course — the students would have produced the learning material for this type of a course.

In other words, students' roles were to: a) bring in resources for the community, b) generate an idea for a project and start a project by informing it to others, c) join an existing project, d) work as a group or as an individual for the project, e) help the community by annotating other people's works. To ensure peer-revision of other students' work, after the ending date students were asked to single out "significant" or otherwise meaningful projects among all the published projects.

Two instructors hosted the course. In practice, the instructors' role was to comment on drafts and final versions of the published projects. The instructors rated the project outcomes individually and negotiated a single grade for each project. The instructors' roles were deliberately diminished in order to ensure community-type feel for the learning environment. When using OurWeb as a platform for learning, it is natural that the flow of communication does not go through the instructor but is distributed, as suggested by Scardamalia and Bereiter (1994).

The students of the course were encouraged to view the activities taken as a form of an online community. The students were referred to as *learning community members*. They were explicitly told that they cannot rely on being taught by the instructors, and they should try to donate their effort and expertise to the community as well as their immediate project groups. The students were encouraged to take advantage of the work of each other rather than re-inventing the wheel (while respecting the copyright). The existing material, projects, and discussions also served as an important source of ideas for identifying new meaningful directions for further exploration and elaboration.

The evaluation of OurWeb concentrates on two issues: 1) time devoted in different activities in the community and 2) community support by openness and transparency of activities in the system.

Community support focuses on issues of joint annotation: is interacting with others straightforward and easy-to-use, and can the discussions be coupled tightly to the context? Yet another issue for community support is the possibility for collaborative process writing: do the community members use it for their benefit?

All of these questions are examined with the statistics logged during the course. Even the times reported below are likely to be fairly accurate due to very extensive logging of client-side events. In addition to the users' navigation, the logs contain e.g. switches of focus from the browser window to another program and back, scrolling of the page, and the display of annotation contents. As a result, a user normally generates several log entries per minute, and even relatively short pauses suggest that the user is not actually working in the system. Considering the time studies, pauses longer than 10 minutes were excluded from the figures below, as well as periods when OurWeb was not the topmost window on the user's display.

3.2 Results

The students spent on average 16.8 hours online. As anticipated, there was a huge variation. Two students spent less than one hour on-line. The maximum time on-line was 65.8 hours. This was exceptional, since the 5 next highest total times were within 24.9-30.7 hours. 8 students were online for less than 10 hours.

The central role of project work shows clearly in the distribution of the time to different activities. The student population as a whole spent only 7.5% of the total time on reading and annotating the background material. Reading and annotating various versions of the project reports (both own and other teams' reports) took 22.2% of the time, and the proportion of writing (in Wiki) was 28.2%. The rest of the time (42.1%) was spent in various miscellaneous ways, in particular browsing the document pool index pages and reading and editing project descriptions.

Individual differences in the distribution of time were small. Only 3 students spent over 15% of their time working with background material, and direct work with the project reports (reading/annotation and writing) in almost all cases occupied roughly half of the total time. 5 people put over 30% of their effort in writing. Most of these had higher than average total time in system.

Distributed processing of vast amount of information can be done by collaborative searching the Web for relevant information and dividing the effort needed to dig out the important points. Highlights can be used in achieving this. The students brought 75 background articles into the joint document pool. In a case where a student brought in the document him/herself, a total of 99 highlights were made to background article. In a case where a background articles was brought into the document pool by other community member, a total of 111 highlights were made. Significant variations existed, from 0 to 78 (average 17.5).

Considering the project work, the students were more eager to highlight other reports than the ones they were contributing themselves: 132 highlights as opposed to 12.

The amount of total highlights (398) and tooltip comments (891) made suggests that it is relatively straightforward to annotate documents. Threaded discussion postings were significantly less, namely 12 postings (of which only one was contributing discussion and 11 were questions about system functionality). Similar results in favour of tooltip comments have been reported earlier (Kurhila et al. 2003a; Kurhila et al. 2003b).

Discussions within a context offered by the joint annotation feature is supported by the figures: there was 164 discussions (comment chains) in a specific location. On the other hand, it seems that the usefulness of the tooltip comments is diminished if the group is smaller; project teams wrote only a total of 28 tooltip comments while preparing their projects. A total of 20 projects were initiated, and 47 draft and final versions were published to the community (average per project 2.14).

4. CONCLUSION

The paper presented OurWeb and discussed its use in a small community-type course. The experience was very positive in several ways, although the size of the student group was too small and work load too large to show all the possible benefits. Even with a small sample, it is evident that OurWeb can open the environment for meaningful collaboration so that each individual can help and benefit from others, by complementing each other's shortcomings as a group (see e.g. Vygotsky 1978; John-Steiner 2000).

Intense interaction over initial ideas and draft versions can transform to knowledge, and the tools offered in OurWeb can appropriately assist in this process. To achieve a high level of expertise there is no need for exceptional intelligence; only long-time participation to expert community and exposure to the knowledge it possesses is needed (Ericsson & Charness 1994).

ACKNOWLEDGEMENTS

This work was supported in part by the Academy of Finland under the Prima and Prose projects.

REFERENCES

Bereiter, C.: Education and Mind in the Knowledge Age. Mahwah, NJ: Lawrence Erlbaum Associates (2002).

Dieberger, A.: Social Navigation in Populated Information Spaces. In A. Munro, K. Hook and D. Benyon (Eds.), Social Navigation of Information Space, pages 35-54. London: Springer (1999).

Ericcson, K.A. and Charness, N.: Expert performance: its structure and acquisition. American Psychologists, 49, 725-747, 1994.

Felder, R. and Brent, R.: Navigating the Bumpy Road to Student-Centered Instruction. College Teaching, 44, 43-47, (1996).

Felder, R. and Brent, R.: Effective Strategies for Cooperative Learning. Journal of Cooperation \& Collaboration in College Teaching, 10(2), 69-75 (2001).

Gardner, A. and Johnson, D.: Teaching personal experience narrative in the elementary and beyond. Flagstaff, AZ: Northern Arizona Writing Project Press (1997).

Grabe, W. and Kaplan, R.B.: Theory and Practice of Writing. London: Longman (1996).

John-Steiner, V.: Creative Collaboration. New York: Oxford University Press (2000).

Kurhila, J. et al.: Peer-to-Peer Learning with Open-Ended Writable Web. In Proc. ITiCSE 2003, pages 173-177. ACM Press (2003a).

Kurhila, J., Miettinen, M., Nokelainen, P., Floréen, P. and Tirri, H.: Joint Annotation and Knowledge Building in Collaborative E-Learning. In Proc. World Conference of E-Learning in Corporate, Government, Healthcare and Higher Education (E-Learn 2003), pages 2249--2252. AACE (2003b).

Leuf, B. and Cunningham W.: The Wiki Way: Quick collaboration on the Web. Addison-Wesley Longmann (2001).

Lonka, K., Linblom-Ylanne, S. and Maury, S.: The effect of study strategies on learning from text. Learning and Instruction, 4, 253-271, 1994.

Miettinen, M., Kurhila, J., Nokelainen, P., Floréen, P. and Tirri, H.: EDUCOSM - Personalized Writable Web for Learning Communities. In Proc. ITCC 2003, pages 37--42. IEEE Computer Society Press (2003).

Munro, A., Hook, K. and Benyon, D.: Footprints in the Snow. In A. Munro, K. Hook and D. Benyon (Eds.), Social Navigation of Information Space, pages 1--14. London: Springer (1999).

Nokelainen, P. et al. Evaluating the Role of a Shared Document-based Annotation Tool in Learner-centered Collaborative Learning. In Proc. ICALT 2003, pages 200--203. IEEE Computer Society Press (2003).

Reed, D. and John, S.: Web Annotator. In Proc SIGCSE 2003. ACM Press (2003).

Scardamalia, M. and Bereiter, C.: Computer Support for Knowledge-Building Communities. The Journal of the Learning Sciences, 3(3), 265-283 (1994).

Suthers, D. and Xu, J.: Kukakuka: An online environment for artifact-centered discourse. In *Proceedings of the Eleventh World Wide Web Conference (WWW 2002)*, pages 472-480.

Tompkins, G.: Teaching and writing: Balancing process and product. Columbus, OH: Merrill Publishing Co. (1990).

Vygotsky, L. S.: Mind in Society. The Development of Higher Psychological Processes. M. Cole et al (Eds.) Cambridge, Mass.: Harvard University Press (1978).

Wenger, E.: Communities of Practice: Learning, meaning and identity. Cambridge: Cambridge University Press (1998).