

RECOLED: A GROUP-AWARE COLLABORATIVE TEXT EDITOR FOR CAPTURING DOCUMENT HISTORY

Masood Masoodian¹, Saturnino Luz², Matt-Mouley Bouamrane² and David King¹

¹*Department of Computer Science, The University of Waikato, Hamilton, New Zealand*

{M.Masoodian, dnk2}@cs.waikato.ac.nz

²*D Department of Computer Science, Trinity College, Dublin, Ireland*

{Matt.Bouamrane, Saturnino.Luz}@cs.tcd.ie

ABSTRACT

This paper presents a usability analysis of RECOLED, a shared document editor which supports recording of audio communication in remote collaborative writing sessions, and transparent monitoring of interactions, such as editing, gesturing and scrolling. The editor has been designed so that the collaboration results in the production of a multimedia *document history* which enriches the final product of the writing activity and can serve as a basis for post-meeting information retrieval. A discussion is presented on how post-meeting processing can highlight the usefulness of such histories in terms of tracking information that would be normally lost in usual collaborative editing settings.

KEYWORDS

Shared editor, awareness, multimedia interaction history and retrieval.

1. INTRODUCTION

Many important texts, such as scientific publications or legal documents, are often created collaboratively. Modes of collaborative writing include the use of a single-user editor asynchronously for revising shared documents over several iterations, and the use of a shared editor synchronously to create the document in one or more collaborative sessions. Although such collaborative activities suffice to produce a final version of a text, they usually do very little to record the history of the writing process. Studies have shown that this type of document history can aid the retrieval of valuable information about the shared document which may not be obvious from the final text itself (Moran et al. 1997, Masoodian and Luz 2001). Yet, support for document history in existing applications continues to be rather limited. Commercial editors have begun to incorporate *tracking functionality*, allowing the software to record data about changes made to documents. Often, however, no recorded information remains across individual sessions about when or in which order each of the changes were made, or how those changes relate to extra-textual content.

Designers of existing groupware applications such as synchronous shared editors have often focused on incorporating interface components which provide means of group awareness, whether it is workspace awareness (e.g. use of tele-pointers, shared scrollbars, etc.) or social awareness (e.g. video links, participant lists or images, etc.). These types of group awareness mechanisms are clearly important in supporting the collaborative work process, and indeed the usability and effectiveness of groupware depends largely on the extent to which it supports group awareness (Dourish and Bellotti 1992, Gutwin and Greenberg 1999).

In most synchronous shared editor software, however, the group awareness information is only used to support the collaborative work activity during the synchronous session itself, after which such information is not recorded and largely ignored. Although some systems record information such as the names of the participants, session duration, and other particulars, the information relating to specific user actions such as text insertion, deletion, focus of attention, and pointing is lost.

In order to address these issues, we have developed a prototype system named RECOLED (REcording COLlaborative EDitor) which not only provides users with important awareness information during the collaborative writing and editing sessions, but also extracts and records extra information (metadata) using the awareness components of the interface for post-meeting access to document history.

2. RECOLED

Collaborative text production typically involves activities other than writing. In particular, communication through face-to-face meetings, or video/audio conferencing often play an important role in the creation of cohesive documents. Speech has been shown to be the most appropriate modality for text revision (Chalfonte et al. 1991) and audio the most effective communication medium for synchronous collaborative work over a shared workspace (Jensen et al. 2000). As such, RECOLED has been designed to be used in conjunction with an audio channel. It is assumed that the co-authors will use the audio channel provided by the conferencing tool for their explicit verbal communication, while the implicit communication and awareness will be supported (and recorded) through the shared editor.

The RECOLED system is based on a client-server architecture, as shown in Fig. 1. These components handle the management of collaborative editing and awareness mechanisms. The server also handles two additional tasks: the recording of audio communication and the capture of the participants' gesturing and editing interactions. The audio tools used by the clients transmit audio in RTP (Real-time Protocol) packets. The audio is recorded by an application which keeps the original RTP timestamps and adds arrival timestamps to the recorded data. These timestamps are vital for detecting speech activities and linking them to text segments. The text management server provides concurrency control between clients through an optimistic locking mechanism to ensure consistency across different copies of the shared document. It is also responsible for the central time-stamping of editing operations. The server uses its copy of the shared document to provide session persistency by sending new clients an up-to-date copy of the document content when they connect to the server. Clients also keep their own copy of the shared document so that the system can have a quick response time to local user events. The clients send their local updates to the server as well as receiving other clients' events from the server.

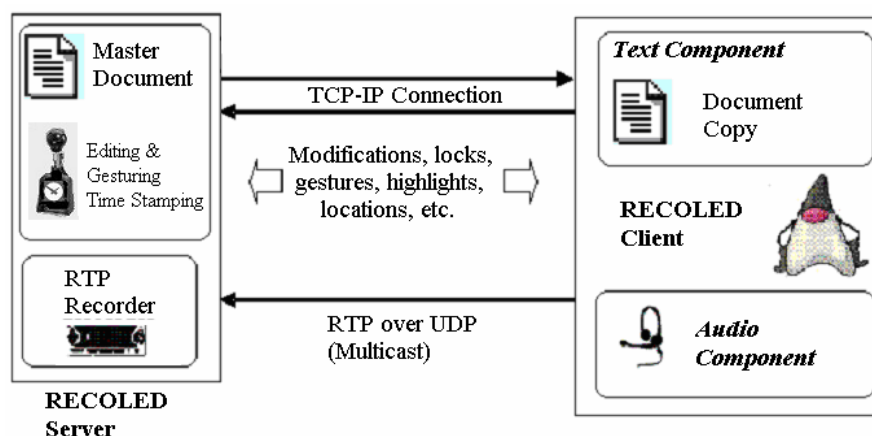


Fig. 1. RECOLED system architecture

2.1 RECOLED Client Application

The aim of the RECOLED client application (Fig. 2) is to provide a multi-user text editing environment similar to conventional single-user word-processing applications. The current version of RECOLED uses optimistic locking to allow users to work on individual paragraphs (i.e. a lock is automatically requested by the system when a user starts editing a paragraph). The lock of a paragraph is also automatically released after the user working on it becomes inactive for a certain period of time (e.g. a few seconds). The system displays a coloured strip to the left of each paragraph to show its lock status (i.e. if a lock has been requested for the paragraph, or if it has been locked and by whom).

RECOLED has a considerable set of awareness mechanisms, some of which are used purely for supporting the synchronous collaborative writing process, while others also provide contextualising information which is recorded by the system as part of the document history. The remainder of this section describes a few of the awareness mechanisms of RECOLED.

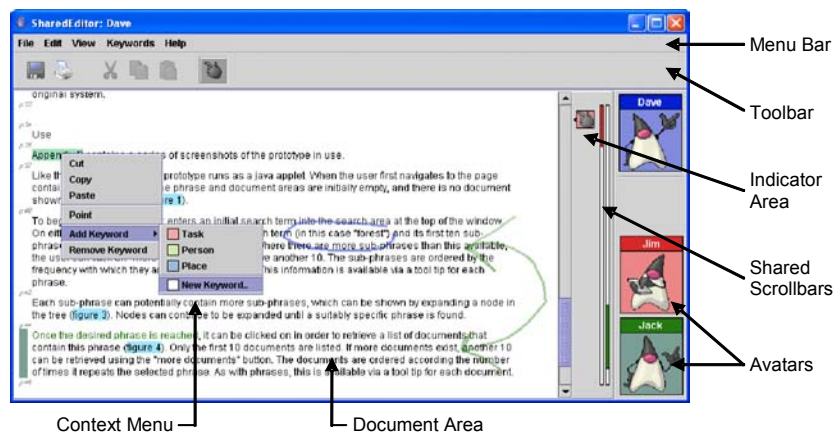


Fig. 2. RECOLED client application interface

2.1.1 Avatars

Each RECOLED client displays a set of avatars to the right of the document area, providing identity and presence awareness by representing all users currently involved in editing the document. These avatars also provide a limited amount of action awareness by taking on different appearances depending on the actions of the user (see Fig. 3). They also serve as a key relating each user to a unique colour, which is then used extensively throughout the interface to identify activities relating to that particular user. The author names are associated with various actions they perform, which are also recorded in document history.




Fig. 3. Client avatars and their actions

2.1.2 Shared scrollbars

Shared scrollbars are shown to the left of the avatars (see Fig. 2). There is a separate scrollbar associated with each of the other current users, displayed in their unique colour. Shared scrollbars provide location awareness by showing where in the document each user is located, and view awareness by showing how much of the document they can see. The users can also choose to synchronize their view of the document with that of another user by linking their views. At present the shared scrollbars are only used for supporting synchronous writing, and their positions are not recorded in document history. This information could easily be recorded should it become useful for post-meeting reference to the document.

2.1.3 Text colouring

Within the document itself, any text that is inserted is initially shown in the colour of the person who typed it (see Fig. 2). Coloured text gradually fades into the normal text colour, which ensures that the document will not become distractingly colourful, while allowing users to gain a sense of the order in which characters were typed.

2.1.4 Gesturing

RECOLED has two different types of gesturing capabilities: simple pointing and freehand drawing. To point at something a user can simply click on that location. This in turn shows the pointed location using a radar pulse, with the centre point surrounded by rapidly expanding circles, which captures and focuses other users' attention on the intended point, as shown in Fig. 4.




Fig. 4. Pointing gesture with radar waves

When the users need to convey more complex information than a simple location, they can hold the mouse button down and freely draw over the surface of the document. Freehand drawings by a user are then displayed on the screen of every other group member (see Fig. 5).




Fig. 5. Freehand gesturing

Both types of gestures are always painted in the unique colour that relates to the user who made the gesture; thus maintaining ownership awareness. Also common to both types of gestures is the fact that they are not persistent and gradually fade away. This reflects the fact that gestures are very temporary and therefore quickly lose their usefulness as the document around them changes and the context of what is said or done at the time of the gesture is forgotten. However, since the gestures are short-lived, they can only be played out to those users who are in the same location within the document as the user performing the gestures. When a gesture is received by a user who is not in a position to see it, an indicator appears to the right of the scrollbar showing the location and owner of the gesture (see Fig. 6). Clicking on this indicator changes the user's view to the location of the gesture, and then the gesture is played back as it was drawn. These indicators, like the gestures themselves, gradually fade over time so that they can be ignored if they are of no interest to a user. If however the user decides to click on this indicator, he or she will be taken to that position in the document and will see the gesture being played. In this case, an anchor is displayed on the user's scrollbar so that it is possible to conveniently return to the original location.




Fig. 6. Gesture indicators

Our experience has shown that it is very common for people collaborating over a document to complement what they are saying through the audio channel by gesturing or pointing over a segment of the document (e.g. saying "look at this section" and then pointing at it). In cases such as these, recording the audio communication without recording the gestures, or at least the location of the gestures, would make parts of the recorded audio communication meaningless.

2.1.5 Highlighting

A common form of communication over a shared document, particularly in the case of asynchronous collaboration, is to mark segments of a document and to leave some form of an indicator, perhaps with a note, for future reference, either for one's own self or someone else. In some systems designed for supporting asynchronous work this is catered for through the use of sticky notes. Although RECOLED is designed primarily to support synchronous writing, it is important that it also supports marking parts of text for future reference. RECOLED therefore includes a highlighting functionality. Highlighting is rather different from gesturing, as it is of a more static nature. To highlight a segment of text the users can select the text and assign a keyword description and its associated colour to the selected text. The keywords can be predefined, or can be defined dynamically by the users, which automatically become available to everyone (see Fig. 7).




Fig. 7. A typical highlighting scenario

3. USABILITY STUDY OF RECOLED

A usability study of RECOLED was conducted to evaluate its capabilities as a synchronous shared editor for supporting collaborative writing. This study involved two separate sets of experiments. The methodology used in both experiments was identical but the tasks that participants had to perform were significantly different. This was done in order to see how different scenarios might impact on the use and perceived strengths or weaknesses of the various features of RECOLED. The aim of this study was not only to evaluate the group awareness and other features of RECOLED, but also to collect meeting data which could be used for post-meeting analysis of the collaborative meeting contents.

3.1 Methodology

The study was conducted at a usability lab. In the first experiment, 14 students took part in seven dyadic writing sessions, and received a book voucher for taking part in the study. The second experiment involved 24 students in 12 dyadic sessions. All the participants were familiar with the use of text editors, but had not previously used a shared editor.

3.1.1 Editing tasks

The group task used for the first study (Task A) required the subjects to collaboratively work on a simple children's story in which the paragraphs had randomly been rearranged and some of its words were replaced with underscores. The participants' task was to cooperatively arrange these fragmented paragraphs in a logical order and to replace the missing words to recreate the story. The task was specifically designed to encourage a high level of interaction and communication between the participants.

Participants in the second study were asked to organize a fictional week-end social function for final year students and staff (Task B). The things they were asked to organize included booking a hotel, travel arrangements, proposing day time and night time activities, assessing costs, etc. Upon starting the session, a basic plan was loaded on to the editor with a number of outstanding issues to be resolved. In order to do this, participants were each handed a (small) set of printed information about the destination's accommodation availability and attractions. Participants had access to both shared and exclusive information. This was specifically designed to encourage interaction and communication between the participants and to ensure that for certain tasks, each person had to rely on the other participant's critical information.

3.1.2 Sessions

Each session began with a demonstration of the software. The participants were then encouraged to explore its functionality and to ask any questions they might encounter. In addition to RECOLED, participants could communicate via an audio channel through the RAT audio conferencing tool, which required the subjects to wear a headset. Once the participants indicated that they were confident in using the program, the experiment proceeded to the actual evaluation task.

There was no time limit set for completing the group task, and the participants were free to continue on the task for as long as they wished. Once the participants felt that they had completed the task satisfactorily, they were asked to fill in a short questionnaire. This was followed by an informal discussion with the researcher, which explored the participants' views about the software, the task, collaborative writing in general, and specific suggestions for improvements to the system. The sessions took about an hour each.

3.1.3 Data recording

The document history capture was done by the automatic editing-gesturing timestamps generation mechanism of RECOLED, which is described in detail in (Bouamrane et al. 2005). Audio recording of the sessions were made through the RTP-recording functionality of the RECOLED server.

3.2 Results and Implications

The use of questionnaires, informal discussions, and general observations made during the sessions helped us gain insight into the participants' view of the software and its features, as well as its overall usability. The main findings of the two studies are discussed below.

Writing. The participants adopted a wide range of writing strategies, which were different not only across the two tasks, but also across different sessions with the same task. This is consistent with the findings of other studies which show that people use a variety of approaches to collaborative writing (Posner and Baecker 1992). In our studies the writing methods also tended to change within each session as participants previously unfamiliar with a co-authoring task, gradually became more comfortable with using the shared editor. An interesting point was to see how quickly participants accepted the idea of sharing a document. As documents were loaded at the start of the sessions, the co-authors made no ownership claims on the document and had no objections in the other participant modifying it in any way.

In the first study the most common strategy was for participants to act as joint writers, working closely together and consulting each other carefully before taking action. This demonstrated the effectiveness of RECOLED as a shared writing tool. Sometimes, however, the participants worked as separate writers, taking responsibility for different tasks or sections of the document and working independently of each other. A common strategy was to have one participant filling in the blank words while the other searched for paragraphs and rearranged them.

In the second study the participants generally worked in tight cooperation, tackling a specific task before moving on to the next one. An interesting point was to see how the various groups perceived priorities differently and therefore chose to address the problems in different orders. Although a basic plan was loaded on to the editor when the session started, the participants did not address the plan items in the order in which they appeared, nor did they try to physically reorder these items within the document. Different groups did however address the items in different orders, according to their own views of priorities (cost, travel arrangements, etc.). It would be impossible to infer these differences by solely looking at the final document. Therefore, the users' perceptions of priorities would have been lost if the planning was recorded using a traditional shared editor. One of our assumptions is that by capturing the time and type of user actions, post-meeting processing of timestamps generated by RECOLED will permit to emphasize different groups' sense of priorities when working on a similar task.

Gesturing. Gesturing was used extensively during the writing process in Task A, although mainly in the form of pointing rather than freehand drawing. Also, gesturing was always accompanied by audio conversations and never on its own. Once again, the fact that some form of gesturing was used extensively, and that it was always in conjunction with audio conversations, demonstrates the value of recording both the gestures and audio in document history, as a way of identifying co-authors' focus during the collaborative writing process.

Gesturing was used very little during Task B. A possible reason for this is that since the planning document was rather short, the participants were always seeing the same thing on their screens. The tight cooperation involved in this scenario, coupled with the availability of the audio channel made the focus of the co-authoring task very clear, thus requiring little or no additional gesturing.

Highlighting. A surprising result was that highlighting was also used frequently during the first study, even though the writing process was synchronous. In the case of this study highlighting was generally used to mark what words were added in place of the missing words. This simple but effective use of highlighting shows that mechanisms such as highlighting are an integral part of co-authoring, and information gained from recording its use can be valuable.

Shared scrollbars. Shared scrollbars were also used heavily during the first study. The participants not only used shared scrollbars for locating the other participant's view, but also on many occasions linked their views to follow one another as they moved through the document. This feature, however, caused some disorientation when one of the participants would for instance forget that their views were linked, and the other participant would scroll through the document.

Text locking. Another interesting result is that the participant almost never edited a single paragraph together during either of the studies. An obvious reason for this is the paragraph-based locking mechanism of RECOLED, which was demonstrated to the users before the sessions began. Therefore, the Participants knew that they could not edit the same paragraph at the same time and never tried to do so. Also, when an audio channel is available, people tend to discuss who should be working on what part of the document verbally. Therefore, the simple optimistic locking mechanism used by RECOLED seems to be sufficient in preserving documents consistency without hindering the co-authoring task.

3.2.1 Participants' feedback

After the completion of the sessions, participants were handed a questionnaire divided into four sets of questions relating to the *task* performed, their view of their *collaboration*, the *usability of the shared editor* and the use of the *audio channel*.

Collaboration. Participants felt comfortable about sharing the document. An interesting point is that when asked to rate their own and their partner's contributions to the overall task, the participants felt that they had contributed equally to the task. They also generally thought that they were rather efficient in completing the required tasks and were satisfied with the final outcome.

Shared editor usability. The participants generally found the awareness widgets to be very useful. As to what could improve the editor, a general feeling was that it would be useful to be able to import other types of objects such as pictures and tables. An encouraging point is that many felt they could easily see themselves using the tool for a number of tasks such as: planning tasks, project and budget proposals, document review, editing papers, computer programming, or as instant messaging or chat tool.

Audio. The ability to speak with each other was usually rated very highly. This is consistent with previous findings (Jensen et al. 2000, Bos et al. 2002) regarding the importance of audio in building trust between participants and encouraging cooperation. As an audio channel was part of the experiment set-up, many participants just took it for granted and assumed that it would not be possible to complete the task without it.

4. INFORMATION RETRIEVAL

A detailed discussion of how the document history captured by RECOLED can be used for post-meeting processing and information retrieval has been discussed elsewhere (Bouamrane et al. 2004). In brief, document histories can support effective retrieval of information from multimedia (audio, text and gestures) recordings of collaborative activities.

A prototype meeting browser called MeetingTree has also been implemented using the concept of temporal neighbourhood, formally defined in (Luz and Masoodian 2005). The final text document of the co-authoring meetings is used as the starting point to the browsing tasks. The editing timestamps captured on each paragraph during the co-writing tasks are used to provide temporal entry points to the meeting audio recording using concurrency of actions and speech segments. These audio segments might in turn point to another set of editing actions performed on different paragraphs (illustrated in Fig. 8a).

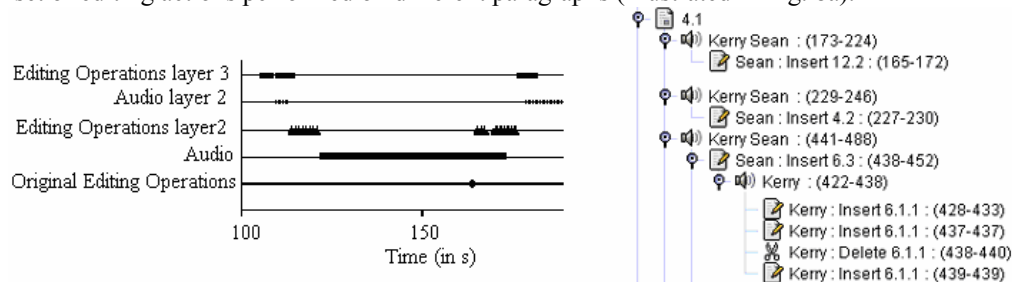


Fig. 8. A Paragraph's retrieval unit and visual representation of retrieval unit

The bottom line shows all operations performed on one particular paragraph. In this case, it is a single atomic operation (gesturing). The line above this shows a concurrent audio segment of participants' speech exchanges. Editing operations layer 2 shows all editing operations performed on *different paragraphs* from the original one, which were performed within the duration of the audio segment. These are then linked to a

different set of audio segments (audio layer 2). Finally, editing operations layer 3 represent all editing operations *not included* in the previous editing layers which happened within the duration of audio layer 2. The retrieval algorithm stops when no more concurrent editing operations or audio segments can be found. A paragraph's retrieval unit then consists of the set of audio speech exchanges and text interactions retrieved by this algorithm. To visualise each paragraph's retrieval unit, we use a tree structure as illustrated in Fig. 8b. This simple technique is able to link each paragraph with several (not necessarily sequential) speech exchanges and other (not necessarily contiguous) paragraphs, while maintaining overall semantic coherence; thereby highlighting semantic and contextual relationships which would otherwise have been lost without the history capture functionality of RECOLED (Bouamrane et al. 2004).

5. CONCLUSION

This paper has presented a real-time collaborative editor that integrates traditional editing and awareness functionality with mechanisms designed to implicitly capture document history. This integration has proven successful in many respects. History capture and awareness feedback mechanisms exhibit a high degree of complementarity, especially in scenarios where communication is performed through speech. Knowledge of what textual objects were pointed at and viewed at different times is as important for post-meeting information retrieval as the main editing operations (e.g. writing and deleting) themselves. From the perspective of the importance of maintaining interaction, the non-linear patterns of editing and gesturing activities on different segments, evident especially in Task B, show that information such as relative priorities and decision rationales do not always remain in the finished product. Therefore, history tracking, used in conjunction with audio, may be the only way to salvage the process by which the text came to be in its final form. Although the importance of text creation histories has been generally recognised in the HCI literature (Moran et al. 97) the work presented in this paper is, to our knowledge, the first study to systematically investigate the significance of this type of information in an interactive framework.

REFERENCES

- Bos, N., Olson, J., Gergle, D., Olson, G., and Wright, Z. 2002. Effects of Four Computer-Mediated Communications Channels on Trust Development. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. Minneapolis, Minnesota, pp. 135-140.
- Bouamrane, M., King, D., Luz, S. and Masoodian, M. 2004. A Framework for Collaborative Writing with Recording and Post-Meeting Retrieval Capabilities. *Special issue on the 6th International Workshop on Collaborative Editing Systems, IEEE Distributed Systems Online*.
- Bouamrane, M., Luz, S., Masoodian M and King, D. 2005. Supporting Remote Collaboration through Structured Activity Logging. *Proceedings of the 4th International Conference on Grid and Cooperative Computing*. Beijing, China.
- Chalfonte, B.L., Fish, R.S. and Kraut, R.E. 1991. Expressive Richness: A Comparison of Speech and Text as Media for Revision. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems: Reaching Through Technology*. New Orleans, Louisiana, pp. 21-26.
- Dourish, P. and Bellotti, V. 1992. Awareness and Coordination in Shared Workspaces, *Proceedings of the Conference on Computer-Supported Cooperative Work*, Toronto, Canada, pp. 107-114.
- Gutwin, C. and Greenberg, S. 1999. The Effects of Workspace Awareness Support on the Usability of Real-Time Distributed Groupware. *ACM Transactions on Computer-Human Interaction*, Vol. 6, No. 3, pp. 243-281.
- Jensen, C., Farnham, S., Drucker, S. and Kollock, P. 2000. The Effect of Communication Modality on Cooperation in Online Environments. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. The Hague, Netherlands, pp. 470-477.
- Luz, S. and Masoodian, M. 2005. A Model for Meeting Content Storage and Retrieval. *Proceedings of the 11th International Conference on Multi-Media Modeling*. Melbourne, Australia, pp. 392-398.
- Masoodian, M. and Luz, S. 2001. COMAP: A Content Mapper for Audio-Mediated Collaborative Writing. *Proceedings of the HCI International Conference*. New Orleans, Louisiana, pp. 208-212.
- Moran, T.P., Palen, L., Harrison, S., Chiu, P., Kimber, D., Minneman, S., van Melle, W. and Zellweger, P. 1997. "I'll Get that off the Audio": A Case Study of Salvaging Multimedia Meeting Records. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. Atlanta, Georgia, pp. 202-209.
- Posner, I.R. and Baecker, R.M. 1992. How People Write Together, *Proceedings of the 25th Annual Hawaii International Conference on System Sciences*, Hawaii, pp. 127-138.