Other Applications

Our emphasis has been to develop technology to support a group working memory. However, the approach is much more flexible than this emphasis suggests. As menioned above, the integration of note taking and video recording is potentially very useful for archival or longer-term activities. Any type of institutional group activity, for example, court-room procedures, would also lend themselves to this type of recording.

It is also possible to use the tool "in reverse" for applications that involve commenting upon existing video or associating notes with real-time video data.

Evaluation Plans

We will evaluate Synthesis more systematically in late 1993 and early 1994. Students taking part in group projects in several courses at Georgia Tech will use it to perform their assignments. These courses include software and user interface design courses in the undergraduate and graduate computer science programs, and a graduate-level collaborative writing course.

In addition, we continue to use Synthesis for our own collaborative work.

ACKNOWLEDGEMENTS

This work was supported in part by NCR Corporation, Human Interface Technology Center. The VISCA Interface layer of Synthesis was implemented by Terry Kane. Kenji Takahashi graciously risked a time-critical writing project on what was then an unproven prototype.

REFERENCES

[1]Dourish, P. and Bellotti, V. (1992) "Awareness and coordination in shared workspaces" *Proc. CSCW'92* (Toronto), ACM, pp. 107-114.

[2] Cooley, T. (1992) *The Norton Guide to Writing*. New York, W.W. Norton, . p 38.

[3] Russell, D., Stefik, M., Pirolli, P. and Card, S. (1993) "The cost structure of sensemaking" *Proc. InterCHI'93*, ACM, pp. 269-276.

[4] Bernstein, M., Bolter, J.D., Joyce, M. & Mylonas, E. (1991), "Architectures for Volatile Hypertext," *Hypertext '91: Third ACM Conference on Hypertext. ACM*, 243-260.

[5] Schon, D.A. (1983) *The Reflective Practitioner: How professionals think in action*, New York: Basic Books.

[6] McMenamin, S. and Palmer, J. (1984) *Essential Systems Analysis*, Yourdon Press.

[7] Posner, I. & Baecker, R. (1992), "How People Write Together," *Proc. 25th Int. Conf. on the System Sciences*, Hawaii, IEEE.

[8] Sharples, M. (1993) "Adding a little structure to collaborative writing" in Diaper, D. and Sanger, C, Eds., *CSCW in Practice: An introduction and case studies*, Springer-Verlag, pp. 51-67.

[9] Sharples, M., Goodlet, J., Beck, E., Wood, C. Easterbrook, S., Plowman, L. and Evans, W., (1991) *A Framework for the Study of Computer- Supported Collaborative Writing*, Cognitive Science Research Paper 190, School of Cognitive and Computing Sciences, Univ. Sussex.

[10] Conklin, E.J. & Yakemovic, K-C.B (1991)., "A Process-oriented approach to design rationale" *Human-Computer Interaction 6*, 357-391.

[11] Grudin, J. (1989) "Why Groupware Applications Fail: Problems in Design and Evaluation," *Office: Technology and People*, 4(3): 245-264.

[12] Badre, A., Hudson, S. & Santos, P. (1993), *An Environment to Support User Interface Evaluation Using Synchronized Video and Event Trace Reporting*, Georgia Institute of Technology, Graphics, Visualization and Usability Center Technical Report 93-16.

[13] Minneman, S.L. & Harrison, S.R. (1993), "Where Were We: Making and using near-synchronous, pre-narrative video" P*roc. ACM Multimedia '93*, (Anaheim, June), ACM, New York, pp. 207-214.

to date have started with existing outlines, and these outline structures drove both the meetings and the review process.

One indication that video provides a significant advantage over audio recordings is the contextual information provided by meeting participants' facial expressions. One note taker decided which parts of the document needed to be written most carefully when he started accidentally noticing puzzled expressions on the faces on the tape. During the meeting he had missed these indications that the document structure was not clear. He also went back and watched more carefully the parts of the discussion that he belatedly realized that he had misunderstood.

Taping also has the unforeseen benefit that non-native speakers are able to review segments of the discussion that they are unable to follow live.

DISCUSSION

Related Work

Posner and Baecker [7] surveyed 22 collaborative writing projects and present a taxonomy of patterns of collaborative writing. They identify four strategies: *single writer*, in which one person writes a document based on discussion with others; *scribe*, in which one collaborator takes down the group's thoughts; *separate writers*, where different people take different parts; and *joint writing*, in which the group writes together. Sharples [8] from a similar analysis [9] identifies three collaborative strategies: the *sequential*, *reciprocal*, and *parallel*.

Synthesis could support any of these strategies with the possible exception of Sharples' parallel strategy—which is most appropriate after the pre-writing stage. The current implementation and our experiences with it, however, are limited to Posner and Baecker's scribe strategy.

We have developed the current version of Synthesis with certain assumptions in mind. For example, we assume that the collaborative group is small (2-5 people), that there is one note taker (this is enforced), that the note-taker is a fully active meeting participant, and that there is one reviewer. We intend to explore alternative assumptions and writing strategies in the future. For example, it is not necessary that the review process be conducted by an individual. A group could review its work collectively, and the video indexing would be a means for each member of the group to recapture points made at the original meeting and make them more precise.

We have integrated our system with a text processing or idea processing tool, Storyspace, which was developed specifically to support the pre-writing or idea organization phases of writing projects.

Using an idea processor to capture group decisions or conclusions has much in common with the practical evaluations of design rationale capture tools, such as gIBIS or itIBIS [10]. Such systems form a promising line of research, and Synthesis is consistent with such use. However, we have set out to support group working memories rather than a permanent archive, because evidence suggests that working memory systems should be more useful and easier to introduce into real work settings. This is because the people who are required to make extra effort to use the technology are those who benefit the most and most immediately. In contrast, groups who record their design rationale can encounter several problems [11]: (1) they may be undiscerning about what they record, and the unobtrusiveness of video would compound this problem; (2) it may be difficult to find relevant information later; (3) one or more group members must do a lot of work to index the archive; (4) those who do the work may not be the ones to benefit later.

The idea of recording group work on video is nothing new. Small group research and protocol analysis studies typically utilize video or audio recordings of work sessions that are analyzed subsequently by a researcher. The idea of transparently indexing a video record by user interface events comes from the I-Observe tool [12] which allows the human factors researcher to collate videotapes of a subject using an interactive application and the keystroke-level events crossing the user interface. The idea behind Synthesis is to turn such techniques round and give them back to users in collaborative settings to support their own work.

A similar proposal is that of Minneman and Harrison [14], whose *Where Were We* allows collaborators in distributed meetings to share videotapes of earlier interactions or dynamic exhibits, such as video segments of a faulty device that the group is trying to fix or redesign. The name of their system comes from one possible scenario of use in which a late-comer to a meeting can rapidly review what was discussed before his or her arrival.

Minneman and Harrison more closely intertwine the note-taking and reviewing activities than we do, thus breaking down the distinction between reflection and action. We too have found the need to record discussions during a review session—something that is not feasible with a single sequential storage medium. However, allowing simultaneous recording and reviewing requires that the meeting participants pay more conscious attention to the presence of the recording equipment. It remains to be seen whether this disturbance affects the process or outcome of collaborative work sessions.

Another difference between our work and Minneman and Harrison's is the relative importance of textual and multimedia information. Synthesis supports the production of collaborative documents. Our video segments are essentially annotations that are attached automatically to textual objects during the note-taking phase. Although video segments may be archived and become valuable artifacts in their own right, we have emphasized their ephemeral role as an external working memory during collaborative writing. In contrast, *Where Were We* is designed as a multimedia production and storage system, in which the textual notes serve as annotations.

information in unpredictable sequences, random-access storage of compressed digital video would be better. Digital storage would also enable us to review and record during the same meetings. Although this is possible using a single tape, it is not feasible in practice, because the seek time for a segment on a two-hour tape may exceed a minute, and the meeting participants are unlikely to wait that long before raising new topics. Another limitation of sequential tape is that cross-tape indexing is impossible. A reviewer cannot therefore review discussions about the same topic that occurred in different meetings, unless the total time of the collection of meetings is less than the length of the tape (typically two hours).

PRELIMINARY OBSERVATIONS

We have used Synthesis for several collaborative projects. Data and preliminary observations from three of these are summarized in Table 1.

	Purpose	Meetings	Participants (note taker asterisked)
P1	Impromptu meeting to plan a series of experiments to evaluate Synthesis	1 (15 mins)	AB, JB*, CP, BM, TK
P2	Refine existing outline of this paper	1 (60 mins)	AB, JB, CP*, BM
P3	Outline an existing paper about a software engineering hypertext tool	2 (60 and 120 mins)	KT*, CP

Table 1: Collaborative writing projects supported using Synthesis

Although we did not plan these projects in advance to set controls, they form an interesting mixture of writing projects. The projects had varying numbers of participants (two, three or five). There were projects with planned outputs and one without, projects about this research and one concerning another research project entirely, projects in which the note taker was the reviewer, and one (P1) in which the note taker and reviewer were different people, meetings that lasted an hour or more and one that lasted just 15 minutes, meetings in which the note taker had prepared an initial outline for discussion (available to the participants in hardcopy as well as on the screen) and one in which there was no formal preparation.

Conduct of Meetings

Most segments fall in a range from 30 to 150 seconds, although in P3, two segments were nearly 20 minutes long. In all meetings, a number of spurious “segments” occurred when the note taker rearranged Storyspace nodes on the screen. Since Synthesis recognizes the selection of a node as a significant topic-related event, even moving a node on the screen would be indexed with a few seconds of video. Most of these segments are easily recognized, because they last less than 10 seconds.

Note takers may play an active role in meetings or may act primarily as the group scribe. In the three projects summarized in Table 1 we find a wide range of degrees of involvement. One acted chiefly as a scribe for the other group members, one participated in the meeting as actively as the other group members, and one “chaired” the meeting by leading the group through the rough outline that he had prepared beforehand.

The note taker plays a pivotal role. All note takers reported that they changed their style of note taking because of their responsibility. Being able to review the videotape later gives the note taker the freedom to take more abbreviated notes than would otherwise be intelligible and to pay more attention to the discussion than to taking notes.

Also, the very presence of a structured note-taking tool induces meeting participants to stick to the point. Participants are more aware of the topic and more likely to stay on it.

Reviewing and Refining the Outline

It seems to be possible to take notes that make sense subsequently (that is, within several days). Other than the spurious segments caused by the note taker rearranging the spatial layout of the outline, most segments are related to the topic that the note taker claimed them to be about. Sometimes, however, fresh insights can be obtained serendipitously when a digression occurs or when the note taker failed to recognize a change of topic during the meeting so that a fragment of discussion is mis-indexed.

We are intrigued by the amount of time it requires for a note taker (or any meeting participant) to become aware of shifts in the topic of the discussion and react to them. A note taker typically does not complete the creation of a new node or revisit an existing node the moment that the conversation changes topic. The current implementation of Synthesis allows a user to review a segment with preambles of various lengths. We are finding that starting the tape about 15 seconds before the start of the segment is the best option. This provides enough context for the reviewer to understand the discussion, without wasting his or her time with irrelevant information. The size of the lag renders precise synchronization of the topic outline with the video record unnecessary.

We expect that the lag time will vary widely between note takers and between writing projects. Our two major projects

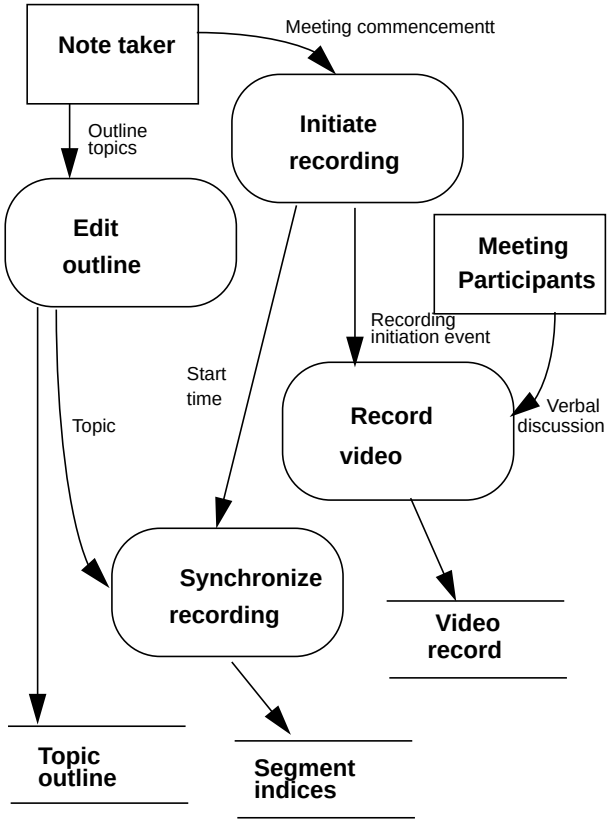

Figure 3: The essential architecture of the recording mode.

video segment to view. To do this, Synthesis must find the point in the video corresponding to that segment of discussion, and play it. In the current implementation, Synthesis does this by retrieving the topic/time association from the segment index file and controlling the VCR through VISCA commands.

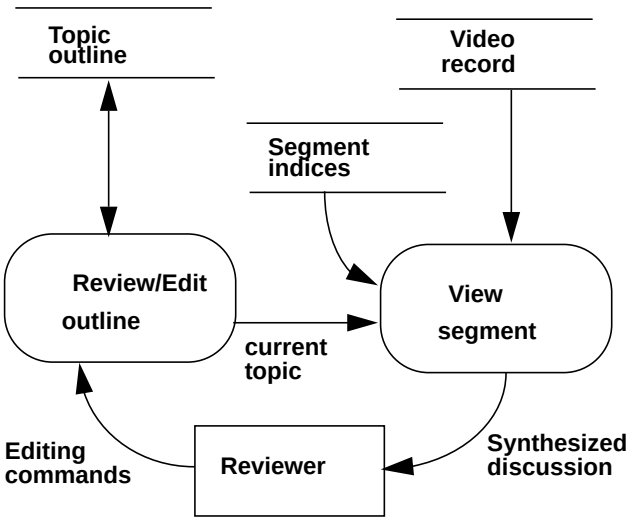

Figure 4: The essential architecture of the reviewing mode.

Current Implementation

The current implementation is an extension of the Storyspace outline processor and hypertext system that runs on the Macintosh. The topic outline is a regular Storyspace outline. Synthesis extends Storyspace by controlling a VCR through the VISCA protocol] and indexing topic-related events (e.g. topic creation or selection) by time of occurrence. The video record is a time-coded analog video recording on 8mm tape (we use the Sony VDeck VISCA-compliant recorder). The segment index is a file maintained by Synthesis that maps Storyspace topic identifiers to time codes. The implementation architecture of the current system is shown in Figure 5

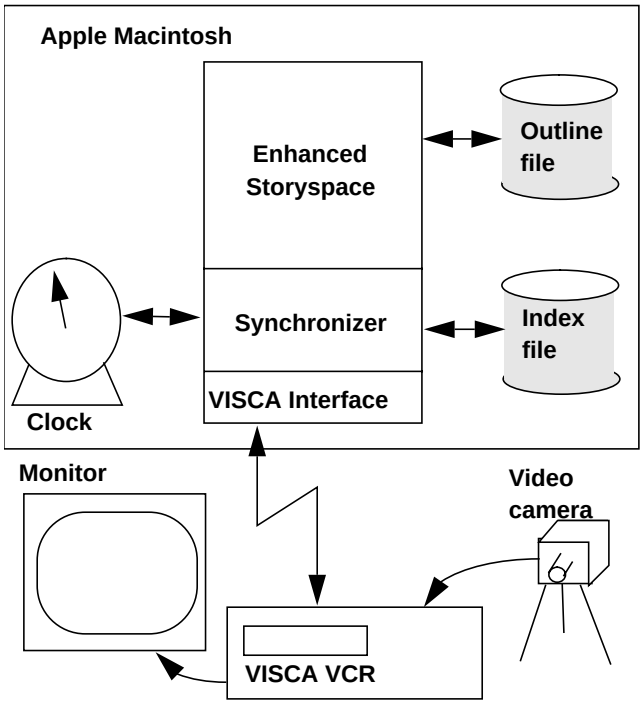

Figure 5: The implementation of Synthesis using VISCA VCR and Storyspace outliner

We wish to evaluate the technology ideas in practical collaborative settings as soon as possible. Thus we have chosen hardware that is readily accessible and easily transportable in preference to state of the art technology.

In the short-term, we have deliberately avoided large installed displays, pen-based input, and networked distributed meeting support. We have instead chosen a standard keyboard and mouse-driven computer as the writing tool and low-end professional video hardware for recording meetings. We only support same-time, same-place interactions, and provide no shared computational support, other than that provided by a computer screen that more than one meeting participant can see.

There are limitations with this strategy. In particular, sequential tape, is not the most powerful technology that meets our needs. Given that users need to browse video

are hats that collaborators wear, not different collaborators. It is possible, for example, to take notes while participating in a meeting; and it is possible to review one's own notes or somebody else's.

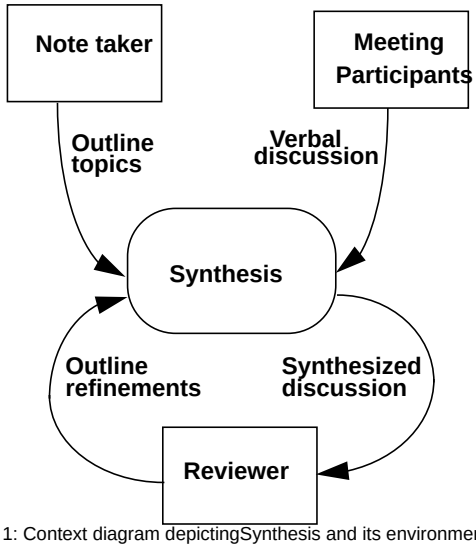

Figure 1: Context diagram depictingSynthesis and its environment.

Figure 2 shows the relationship between the two modes of use: recording and reviewing. The two processes that mediate these modes produce and consume information of three types: topic outlines, video records, and segment indices. The topic outline is the current version of the group's document. The video record is a video and audio recording of the meeting. The segment indices associate topics in the outline with contiguous segments of video.

Recording Mode

Recording operates when the participants are immersed in a discussion. Only the note taker needs to reflect on the structure of the discussion while it is being recorded, and the transparent indexing performed by Synthesis means that the note taker does not have to index the tape manually or strive to keep complete notes.

In Figure 3, the recording task is decomposed into four logical processes: edit outline, initiate recording, synchronize recording, and record video.

Edit outline. Outline editing could be done using any outline processor, word processor, or design or planning tool. The current prototype uses Storyspace.

Initiate recording. The note taker is responsible for initiating, pausing, and resuming recording during a meeting, and terminating the recording at the end of the meeting. The prototype version controls the VDeck directly whenever the note taker invokes the VCR control options from within Storyspace.

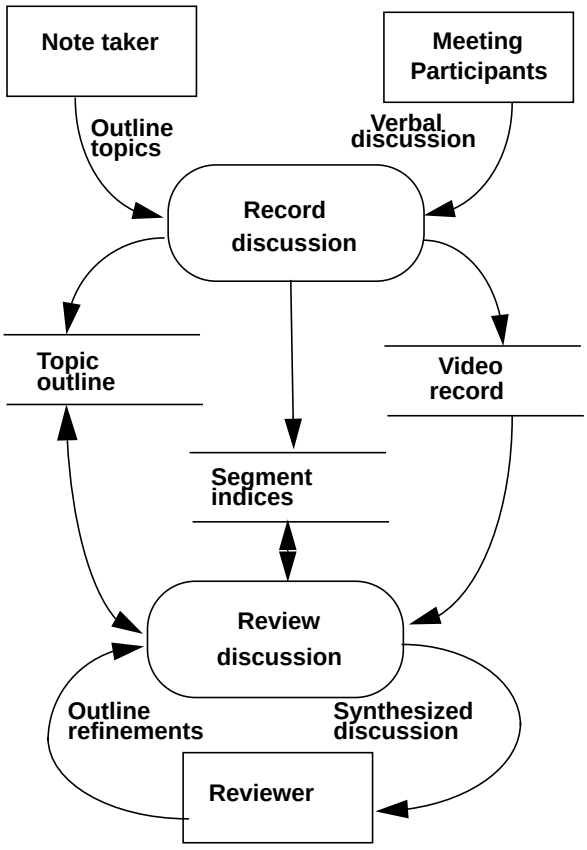

Figure 2: The two modes of use and the information they share.

Record video. The system somehow records the meeting onto a permanent medium. The current implementation uses a standard video camera and a VISCA-compliant VCR under the control of the Synthesis software.

Synchronize recording. Synthesis associates any topic-related event with the topic it is about and the time offset from the start of the recording. In the current implementation, a segment is defined as any contiguous segment of tape starting at the time of a topic-related event and ending at the time of the next one or the end of the recording.

Review Mode

Reviewing, in contrast to recording, is a more reflective activity, in which the reviewer must pay attention not only to the content of the discussion but also its structure. The reviewing process is shown in Figure 4. There are two main sub-tasks: Review/edit outline, and View associated segment.

Review/edit outline. Here the reviewer selects topics to visit, reads the associated text or enclosed topics, and modifies it on the basis of his or her memory of the discussion. In the current implementation, this function is supported directly by Storyspace.

View associated segment. If the reviewer wishes to see a segment of video from the meeting to jog his or her memory, the reviewer may select a topic and call up the associated

subject matter; (2) it is used by the group during the performance of some task (e.g. writing a document), as opposed to being recorded for future purposes; (3) it is of comparatively short duration (in our case the interval between successive meetings, or at most the duration of an entire writing project), as opposed to being a permanent archive; (4) some of its contents are re-encoded into a semantic representation in long-term memory (some of the recorded discussion is distilled or summarized during the writing process); (5) rehearsal, the reactivation or replaying of information in working memory, facilitates its transfer to long-term memory (selective reviewing of videotaped discussion segments facilitates the production of a draft.)

Structure of this Paper

In the next section we describe *Synthesis*, an experimental video-based group working memory tool for collaborative pre-writing. We emphasize its *functional architecture* over its current implementation, and we describe how each of the functional components is currently implemented.

We then present some preliminary observations from three writing projects that have used Synthesis (one of which produced this paper). These observations are not intended to serve as an empirical validation of the group working memory concept, but do suggest some priorities for more controlled studies.

In Section 4, we discuss our preliminary results, and we contrast the mechanisms implemented in Synthesis with the policies of use that teams might prefer for a range of collaborative writing projects. We also discuss related work and some other applications, not involving writing, and not necessarily collaborative.

SYNTHESIS

Synthesis is an experimental video-based group working memory tool for teams. It enables team members to recall the details of verbal discussions that would otherwise be forgotten.

Meetings are recorded on videotape while a note-taker records or edits an outline representation of the document in production using the *Storyspace* outline processor [4]. As the note-taker edits or creates Storyspace topics, Synthesis indices each with the current video segment (actually, the topic creation or visitation time and the duration of the segment until the next topic). Subsequently, one of the meeting participants (the "reviewer") reviews the outline produced during the meeting. Using this information, the reviewer elaborates the notes taken during the meeting.

The two modes of use that Synthesis provides, recording and reviewing, correspond directly to the immersive and reflective forms of engagement that skilled practitioners engage in [5]. The recording mode requires little extra effort from the users; the meeting is being recorded continuously unless the note-taker intercedes to switch recording off. While it is recording, Synthesis indices the discussion transparently. This invisibility corresponds with the unreflective nature of most discussions.

The reviewing or sense-making task, however, is inherently more reflective, and it is here that the technology is more visible. The reviewer uses the outline recorded in StorySpace and the indexed videotape reflectively and with deliberation.

The Experience of Using Synthesis to Write

In a typical Synthesis session, the camera is set up in the meeting room to capture the faces of the participants. The notetaker turns on the camera and VCR as the meeting begins. The notetaker then uses Synthesis to build an outline of discussion topics. As one of the participants brings up a new point, the notetaker creates a new topic marker in the Synthesis outline. Whenever a new topic is created, Synthesis interrogates the VCR to learn the tape location. Synthesis then associates that tape location with that topic. Also, whenever the participants return to a topic that has already been discussed, Synthesis associates the current tape location with that topic. At the end of a meeting, therefore, Synthesis has generated a list of topics and an associated list of videotape locations where each topic has been discussed.

Later, one of the participants (usually, but not necessarily the notetaker) refines the notes taken during the meeting. Whenever something in the notes is not clear and the reviewer wants to be reminded exactly what was discussed, he or she instructs Synthesis to play the video segments associated with the current topic. If there are several segments, the writer is given a choice.

Essential Architecture

The essential architecture of Synthesis is shown in Figures 1 to 4. The term "essential architecture" [6] refers to the functional architecture of a system shorne of all implementation restrictions. Thus the essential architecture depicted in Figures 1 to 4, is the functional architecture for any group working memory system of the type we are investigating. Different systems, of course, have different implementation architectures; that is, their essential functions may be maped onto completely different devices and subsystems. Where one system might permit the interleaving of two functions, another might force them to be perfomed sequentially. The current prototype of Synthesis is just one possible implementation architecture.

As is customary in information systems analysis, the essential architecture is depicted in a series of data flow diagrams. The arcs in these diagrams represent abstract interfaces, not sequences of events. Rectangles represent user roles; rounded rectangles represent functions; parallel lines represent components of the group working memory.Figure 1 is a context diagram showing the interfaces between Synthesis and the people in its external environment. People play three roles with respect to Synthesis: one person takes notes during a meeting (note taker), one reviews the discussion and notes subsequently, and a group of collaborators contribute to the meeting verbally (meeting participants). These roles

COLLABORATIVE PRE-WRITING WITH A VIDEO-BASED GROUP WORKING MEMORY

Colin Potts
College of Computing, Georgia Institute of Technology
Jay David Bolter
School of Literature, Communication and Culture, Georgia Institute of Technology
Albert Badre
College of Computing, Georgia Institute of Technology
{potts@cc,jay.bolter@lcc,badre@cc}.gatech.edu

KEYWORDS

Computer-supported cooperative work, Collaborative writing.

ABSTRACT

Synthesis is a computer-controlled multimedia tool to enhance group communication during the early stages of collaborative work. Synthesis operates in two modes: one for recording collaborative sessions and one for playback and editing. A face-to-face meeting is recorded on videotape while one member of the group takes notes using Storyspace, an outline processor. Synthesis automatically associates the current video segment with that entry in the outline. Later one or more of the participants can review the notes and call up the appropriate video segment. Synthesis provides an external working memory for information that would otherwise go unrecorded and forgotten. Our early experiences using it for real writing projects suggests that the video record is used in a variety of ways to jog the writer's memory for discussion elements.

INTRODUCTION

Collaborative Writing and Pre-writing

Collaborative writing is a form of intellectual teamwork that yields written text. Since most non-routine collaborative projects produce texts—such as plans, designs, proposals and recommendations—the definition of collaborative writing is potentially very broad.

Collaborative writing research has emphasized how important it is that collaborators understand what other members of the team are doing and how they coordinate their work [1]. In this paper, we consider the *prewriting* phases of collaborative writing during which the collaborators discuss the ideas a document should contain and how it should be organized. Prewriting is a common term to describe the idea generation and sometimes the organizing stages of the writing process [2]. These tasks contrast with later phases of the writing process, by which time the collaborators have agreed on what the document should say, and work to articulate their message in the most effective style.

Pre-writing activities are typical of many collaborative tasks that are not usually thought of as writing tasks.

Sense-making and Group Working Memories

When professionals collaborate on intellectually demanding projects, they typically spend only a small portion of their time working face to face. Much collaborative work is done by individuals between meetings in pursuit of collaborative goals. When the meetings are held to produce a planned artifact, such as a proposal, it is important that the notes taken during the meeting in some way contribute to the eventual output of the project.

There are two problems to overcome, however. First, written notes do not capture all of the textual material (written and spoken) that is produced in the meeting. In fact there is tremendous compression of that material in most written notes. Second, written notes do not capture the subtle paralinguistic information that can indicate closure or skepticism.

Collaboration therefore often takes the form of *sense-making* episodes [3], in which individual collaborators or the group as a whole tries to remember where it got to last time and impose some organization on scattered ideas and notes.

We are investigating group use of computer-based external "working memories" to facilitate collaborative writing tasks.

By analogy with human memory, an external "group working memory" has the following characteristics: (1) it uses a superficial encoding (video of verbal discussions in our case) as opposed to a rich, semantic representation of the