

WorkspaceNavigator: Tools for Capture, Recall and Reuse using Spatial Cues in an Interactive Workspace

Arna Ionescu, Maureen Stone*, Terry Winograd

Stanford University, USA

{arna, winograd}@stanford.edu

*StoneSoup Consulting, USA

stone@stonesc.com

Abstract. This paper describes the WorkspaceNavigator, a suite of tools to support the capture, recall and reuse of material and ideas generated by a group working in an interactive workspace in an unstructured manner. Our focus is on the capture of digital information, including screenshots, files, and URLs. These are stored as a sequence of timeslices, integrated by an overview image of the physical space at the time of the capture. The overview image provides spatial cues for accessing the captured information. We conducted two user studies of the WorkspaceNavigator tools and present a number of observations on the design of the interaction and the ways in which users understood and adapted our tools. We show that capturing coordinated slices of digital information is useful for recall and summarization activities, and that coordinating access through the visual metaphor of the overview image is understandable and effective.

Introduction

Interactive workspaces are technology-rich environments for displaying and manipulating information of many forms on interlinked shared and private computers. A typical group working in such an environment may collaboratively browse the web, make presentations, draw in a digital notebook, construct lists, schedules, and documents, or work with other programs. All these activities can be performed either on the shared displays or on users' private laptops. It is easy to lose track of information during such a meeting, and nearly impossible to

reconstruct it afterwards from individual file systems and browser history lists. Even people who were in the meeting generally don't remember all of the things that were relevant and it is even more difficult for people who were not there but need to understand what happened.

This paper addresses the problem of recalling information from a session either during the meeting or some time later. To date, our experience is that recall is currently used mainly for creating meeting summaries. Our hope is that future work will demonstrate that people will use the information for other purposes, such as memory augmentation, as they become more familiar and comfortable with new tools.

Our primary interest is in material generated incidentally as people work together naturally, not as a result of people consciously organizing their information for retrieval. We have developed a collection of tools called the WorkspaceNavigator to capture unstructured work, to help users quickly find and reuse previous materials, and to help users recall where they were in their work by reviewing a timeline of overview images. We use a combination of screen snapshots, file lists, and digital camera images to capture activity during a meeting. With the WorkspaceNavigator, users can go back in time to recall what information was displayed on the various screens in the room, and can easily reopen files and applications that were previously in use. This automatically captured information can be augmented by user-supplied annotations both during and after a session.

A variety of meeting capture tools have been developed (see section on Related Work) using audio or video recordings as a basis for enabling review and recall of meeting activities. In general, these are focused on the words and actions of the people, rather than on the electronic material they are manipulating. We chose to capture discrete content such as the computer displays and participant locations at distinct points in time. We take advantage of people's spatial memory by creating an overview image of the room and its participants, which users can click on to delve deeper into the information, accessing screenshots and files from the various displays in the space. Thus, spatial and visual browsing is designed to be the primary access path to the information we capture, augmenting basic search and file browsing.

We chose not to use video or audio capture, seeing whether we could achieve the value of a video record in a more accessible, less storage-intensive snapshot form. Indexing into audio or video from other cues, such as the form of drawings or content on slides (Abowd, 1998; Minneman, 1995; Moran, 1997; Pederson, 1993; Whittaker 1994) has already been demonstrated.

For our research, we have employed the iRoom (Johanson, 2002b), a prototype interactive workspace that includes three wall-mounted displays with touch screens, a tabletop bottom-projected display, video cameras, and a wireless network that enables the use of laptops and PDAs as an integral part of the collaboration. It also contains a variety of wireless input devices we call iStuff, such as buttons, sliders, knobs, and even a toy dog that reacts to being turned upside down (Ballagas, 1993).

The iRoom software is based on the iROS middleware system (Johanson 2000a), which supports a number of applications for sharing and moving

information. Multibrowse (Johanson, 2001) allows users on any computer to move content between browsers and other applications on the shared screens and on laptops that enable this capacity. PointRight (Johanson, 2002c) allows any pointing device in the room to control a cursor on any of the shared screens. These, along with more conventional facilities such as shared network disk drives, provide a flexible, integrated space in which people can work together.

In our efforts to develop an effective suite of tools for the WorkspaceNavigator we have tested the prototype with two groups of subjects cooperatively working in the iRoom. One was a group of business educators and students conducting a session of their own work. The other was an experiment with six young adults in three sessions, performing a task of our design. Results from both groups provided us with insights as to which specific tools in the WorkspaceNavigator are useful and which metaphors and techniques work. These studies have suggested further extensions and improvements, which are being incorporated in future applications of the WorkspaceNavigator.

Related Work

Meeting capture, including the infrastructure required and the interfaces for presentation of material before, during and after a meeting, has been a popular topic for research. Our focus is on retrieval and access methods in order to assist in recall and reuse of material generated during informal meetings in technology-rich environments such as an interactive workspace.

Researchers have looked at different methods of displaying captured information, depending on what information they chose to capture. The most popular dimension to use when displaying captured information is that of time. eClass (formerly Classroom 2000) (Abowd, 1998), NoteLook (Chiu, 1999), LiteMinutes (Chiu, 2001), Designer's Outpost (Klemmer, 2002), Coral (Minneman, 1995), Time Machine (Rekimoto, 1999), and TeamSpace (Richter, 2001) all rely on time to help users navigate through material. The timelines are presented in a variety of different forms, including web pages, email, and custom applications. While Designer's Outpost and Time Machine both use images of the electronic space (screen shots), none of these projects uses images of the physical space to help guide users' browsing. EClass, LiteMinutes, NoteLook and TeamSpace all provide links to the video from particular moments, but the video images are not used as indexing cues for finding information.

Other projects focus on organizing materials by creating links between information, either automatically or with user definitions. NotePals (David, 1999) and Dolphin (Mark, 1997) both create repositories of linked information, but without a timeline. Taking links a step further, Tivoli (Pederson, 1993), eClass (Abowd, 1998) and FiloChat (Whittaker, 1994) use physically created objects, such as timestamps or pen strokes, to index audio streams. Minneman's (1995) Coral project synthesized continuous multimedia streams of data from different sources and provided user-created and automatically created indices into the information captured. While there are many conclusions in common between our studies and Minneman's (the value of annotations, the usefulness of such a

system for creating summaries), they use a significantly different capture model (continuous audio and video). They also do not explore using a snapshot view of the meeting as part of the access and recall mechanism. Moran (1997) performed a now classic study utilizing the Coral system showing how users come to rely on the linkage of audio with pen strokes when creating meeting summaries if such a mechanism is available. Our work suggests that linking overview and screen snapshots with file data is similarly useful in an environment that encourages multiple sources of electronic data.

Other systems focus on organizing information for recall. Czerwinski (2002) looked at users' recall of activities completed on their desktops both 24 hours and 1 month after a work session. She used both condensed video and still images taken during the work session to help jog users' memories about their activities. Other systems automatically create a text or icon-based history of a user's actions, which can later be searched using various parameters (DuVaul, 2000; Lamming, 1994; Rhodes, 1996).

Our hypothesis is that the spatial arrangement of an interactive workspace is a strong basis for retrieval of the electronic contents. The strength of human spatial memory has been harnessed by a number of researchers in designing interfaces (Fass 2002, Robertson, 1998). To give a concrete example, Renaud (2000) showed how visual and pictorial representations are far superior to verbal descriptions when trying to resume work after an interruption.

Our research focuses on capturing work performed in interactive workspaces, and our approach diverges from previous systems in three key points. First, we capture information using a snapshot model, rather than a streaming audio and video model. Second, we rely upon a spatial organization to present the material to users. Third, we provide multiple layers of access to the information so as to help users better navigate and quickly reuse past materials.

WorkspaceNavigator

The WorkspaceNavigator is a collection of coordinated tools that together create a picture of what occurred in an interactive space. Each tool captures a different kind of information, and can be applied depending on needs and availability of infrastructure.

Our theoretical framework begins with the concept of a *virtual workspace*, which is an abstraction of a body of work (typically a single project) that will build a collection of shared information over time.

Each virtual workspace in turn contains a series of *sessions*, each taking place for a bounded period of time. A group meeting is a typical session, or users may interpret a single meeting as a sequence of distinct sessions if the topic changes. Each session takes place in a physical location, or *setting*. There can be different settings for sessions in a single virtual workspace, e.g., when a project holds meetings in more than one place. As a result, material captured in one setting may be recalled and reused in a different one. We provide an interface for users to declare the start, end, and identity of sessions and settings, making no attempt to infer this automatically.

At regular intervals during a session (our current default is 30 seconds), we capture and store a *timeslice* containing a snapshot of the information in the physical space. Each tool captures information upon the receipt of a centrally deployed synchronization event. Annotations can be added to timeslices to aid retrieval or simply to add relevant information. The captured information is stored in a network-accessible file directory with meta-information (time stamps, annotations, etc.) in a PostgreSQL database.

Capture

At every timeslice, capture tools in the iRoom can gather the following information:

- Wide-angle overview image of the iRoom
- Screenshots of every enabled device in the iRoom
- URLs open on all enabled devices
- Application files (Microsoft Word, Excel and PowerPoint) open on all enabled devices

For both the URLs and the files, only the name is captured. We do not copy or otherwise preserve the contents of the file or URL at the time of the timeslice. While this would be possible to do for files, its extension to URLs, which point into the ever-changing web, is more difficult.

A machine is “enabled” for capture when it is running our capture application, written in Java 1.4 with Python extensions. This is an executable that is always running in the background of the permanent displays in the iRoom, and that users can choose to start on their laptops when joining into a session.

Since we are capturing information from private laptops, we give users easy control over stopping and starting recording by providing a two-state “pause” and “resume” button in the lower right corner of a device running the capture software. If users are working on something personal or sensitive during a session (such as email), the machine can still participate in the workspace without being subject to unwanted recording.

Annotations

We provide a number of tools for users to make annotations. First, when a new session is initiated, the user is prompted for an annotation that will be associated with the initial timeslice for that session. Second, while users are reviewing information from a session, they can enter annotations into the text box on the WorkspaceNavigator for the currently shown timeslice.

Our initial hypothesis was that people would not want to stop and make annotations during a session, as it would interrupt the interactive flow. However, initial testing showed that users wanted to make annotations in real-time, so we provided two means to do this. The first is an application tool that runs on any machine (usually people’s laptops), which supplies a simple text box into which an annotation can be typed. When an annotation is submitted, it is attached to the nearest timeslice. The second tool, which we designed to be less obtrusive, was a

simple mechanical iButton (component of iStuff (Ballagas, 2003)) that a user can press at any time during a session. At the beginning of a session, the button is associated with a specific user, and whenever it is pressed it adds the annotation “*USERNAME* pressed the iButton” to the nearest timeslice.

Annotations are public; any user who browses the information will be able to see all annotations. We hypothesize that users will benefit from seeing annotations made by all participants, rather than just their own.

Information Interface




Figure 1. The WorkspaceNavigator interface contains a timeline of the sessions within a workspace, a snapshot window for the wide-angle camera image of the selected timeslice, and a text area for annotations.

Our prototype interface, written using Java 1.4, includes a number of dialogs for users to log into a virtual workspace, to identify, start, and end sessions, and to display layers of captured information. It was built as a quick way to experiment with different mechanisms and will be the starting point for further design. The main interface is shown in Figure 1.

The focal elements of the interface are a timeline displaying sessions and an overview wide-angle image taken during the currently displayed timeslice. Users can navigate through the captured timeslices either by directly clicking on the timeline, or by using the arrows to step forward and backward. The text area near the bottom of the screen displays any annotations associated with the timeslice.

Users move deeper into the information by clicking on the image of a display in the overview image. This opens, in a separate WorkspaceNavigator window, a condensed screenshot with an active listing of files and URLs open on that device when the shot was taken, as shown in Figure 2.

Any of the files or URLs captured, as well as a full-resolution version of the screenshot, can be reopened and displayed on any device currently available in the iRoom. Future plans include removing the active file list and having users open files and URLs by simply clicking on their representations in the screenshots.




Figure 2. A screenshot window shows the captured screen image, along with a list of files and URLs associated with it.

USER STUDIES

We conducted two qualitative user studies to evaluate the effectiveness of parts of our suite of WorkspaceNavigator tools.

Study 1: Entrepreneurship Meeting

Subjects

The subjects for this study were a research group from a business school. They approached us requesting the use of the iRoom for an intensive three-day meeting as part of their study of the entrepreneurial process, and were willing to be test subjects for the WorkspaceNavigator. The group consisted of academic researchers, students, and industry consultants. This group was relatively technologically unsophisticated, with a limited understanding of the underlying technical model of the iRoom and the WorkspaceNavigator. However, they were very interested in using our technologies to capture every possible detail of their sequence of highly structured meetings. The number of actual users in the iRoom at any given point ranged from 5 to 13.

Setup

The subjects worked with the 3 wall-mounted SMART Boards (www.smarttech.com) and 5 laptops connected wirelessly to the network and running the iRoom Software. The laptops were often projected onto the SMART Boards using VNC (Virtual Network Computing, www.uk.research.att.com/vnc/). They also made occasional use of a paper flip chart on a moveable easel. Two video cameras were mounted in the iRoom for their own research purposes.

Each computer in the room was running the capture software and had the WorkspaceNavigator installed. Laptops also had the text annotation software installed. Each day, three users were given iButtons configured with their names.

Protocol

The main purpose of this study was to validate our capture model and to see how the WorkspaceNavigator provided useful tools. This study also created a database of information with which to refine our implementation of the WorkspaceNavigator.

This study contained three distinct parts:

- A preliminary meeting, several months before the main meeting, to familiarize the participants with the iRoom and the WorkspaceNavigator. As a result of this preliminary study, we made substantial changes to our implementation, adding the real-time annotation facility, the overview images to augment the screen capture and file listings, and the privacy button.
- The three-day meeting, which included participants from the preliminary study as well as new people.

- A follow-up meeting a week after the main study with the researcher group only, to help them use the WorkspaceNavigator for review and analysis.

At the end of the 3-day session, all users were asked to fill out short, mostly qualitative questionnaires about their experiences in the iRoom, and were interviewed as a group. Private interviews with selected participants followed. For the preliminary and follow-up sessions, we took notes and performed a group interview.

Observations

These users were very excited about the captured information the Workspace-Navigator tools provided. They made extensive use of the screenshots and made a large number of real-time annotations as the meetings progressed. They did not spend significant amounts of time in the meeting navigating material using our information interface, but, with our assistance, did use the interface afterwards for review. Every evening the users created summaries of the day's activities using the screenshots, which they then sent out to all involved. They also used the screenshot summaries to orient participants who had missed the previous meeting. We found it interesting that in writing summaries they browsed for the screenshots rather than using the files they created or their own video recordings of the meeting. This suggests that the screenshot record is a powerful and manageable way to capture and recall meeting activities.

The subjects made extensive use of both text annotation and the wireless iButtons. During the second day of meetings, the longest day, they made 147 text and 45 iButton real-time annotations. About 40 timeslices had multiple real-time annotations, either from the same person or different people; when something particularly exciting happened, users often became quite zealous about marking the spot in the session. The content of annotations differed according to who wrote them. As can be seen in Table I, some users wrote out full sentences, others abbreviated as much as possible, and some seemed to have special codes such as "xxxxxx." When asked in the questionnaires why they made annotations, one user responded, "[I] wanted to use snippets to create a video and had annotated with that in mind." It is clear to us that the textual content of the annotation was not necessarily important, as illustrated by the use of the iButtons and by the fact that they often typed no content into the text annotation application before pressing submit. Most important seemed to be marking the point in time, with content to be extracted later.

During our initial pilot session, we observed that explicitly starting and ending sessions helped organize the flow of the meeting. Users had to stop and review what they were discussing and explicitly change topics, which helped get all participants on the same footing. However, during the three days of meetings, users often forgot to start or end sessions. The most dramatic example was the session they left running overnight after the second day, suggesting that some automatic session management would be helpful.

Timeslice 1	Jane pressed the iButton Jane pressed the iButton Jane pressed the iButton Jane pressed the iButton Robert pressed the iButton What we need is..... Door to that is Cell phone & xml and 10 million usersxxxxx – Jane
Timeslice 2	-- Jane close this session Carl -- Jane -- Jane -- Jane
Timeslice 3	15 minutes into the vision session, Sean has drawn upon his own website plus made a preliminary schema to explain his narrative structure, and also drawn Susan's narrative on smartboard one. Would be nice to know when he moved to smartboard one or back to smartboard 3. -Robert At 17:40, the Vision Framework was projected onto the middle screen. This framework was used to help elicit the vision from Susan. How helpful was it? – Carl

Table I. A sample of real-time annotations made to timeslices. User names have been changed.

People using laptops did not use the pause/resume button, in spite of their stated concerns about privacy. While we explained each morning how the button worked, the users not only forgot to use it, they forgot it was available. Of the 13 users that filled out our questionnaires, only two people noticed the existence of the buttons, and neither ranked them as being very useful. When these two users were asked about privacy concerns, both commented that they did have concerns. One even wrote, “Yes [I do have privacy concerns], we need to be able to delete some images from the drives because some content is proprietary.” This suggests that a way to remove sensitive information afterwards may be more effective than a real-time privacy mode.

Study 2: Planning a Trip

The previous study taught us a great deal about our capture model and its application, but little about the use of the information interface. Therefore, we designed a study in which using the information interface for review was part of the task.

Subjects

Six technologically savvy young adults, evenly split between males and females, about to enter their freshman years at college participated. They were told that they were participating in a study looking at how they use the iRoom, particularly at how they recover previously referenced material. They received no monetary compensation for their participation.

Setup

The subjects used the three wall-mounted SMART Boards and one laptop on the network running the iRoom software. A video camera in the back corner of

the room recorded sessions for our research purposes. Data, such as the websites viewed, were automatically captured by the WorkspaceNavigator tools, while data pertaining to human actions, such as how many times the information interface was consulted, were logged by an experimental observer sitting in the room.

Protocol

The six subjects were split evenly and randomly into two groups. Group 1 visited the iRoom three times over the course of a week, with Group 2 joining them during the final visit. Before starting the study, both groups were given 45 minutes to learn how to use the iRoom technologies; we gave a short explanation of the hardware and software, and then three short warm-up tasks to get the users comfortable with the technology. The activities for the three sessions were:

- Session 1, Day 1 (Group 1): Subjects were instructed to spend 1 hour planning a 3-day trip for themselves to any place they desired (they chose Bangkok). They were given a total budget of \$2000 and were instructed to keep an itinerary using Microsoft Word, a budget using Microsoft Excel, and to run the WorkspaceNavigator to capture their work.
- Session 2, Day 4 (Group 1): Subjects were instructed to spend 1 hour planning a trip to the same location as they chose in Visit 1. However, they now had a total budget of \$5000. They were encouraged to reuse information from Session 1.
- Session 3, Day 8 (Groups 1 and 2): Subjects were instructed to plan a 2-week trip for all six subjects to the previously chosen location. Their total budget was \$8000, and they were to use the same tools as in previous sessions.

These tasks were explicitly designed to emphasize the use of recalled material. Although they are not typical of all meeting tasks, they are representative of a common class of tasks for which recall tools can be of special value.

At the end of the each session, subjects were jointly interviewed about their reactions to the iRoom. At the end of the final visit, all subjects filled out a short questionnaire.

Observations

As expected, users relied upon the information interface increasingly during each session. In session 2, once users had a history available from session 1, they opened six previous documents using the interface and in session 3 they opened twelve previously used files or URLs. Sometimes the users relied on the keyword search to find the captured material they were looking for, and other times the users found information following the visual cues provided by the interface. In Session 2, for instance, we observed a subject looking for a nature tour that had initially been identified in Session 1 but was too expensive for that session's budget. Video analysis revealed that in less than 3 minutes, including distractions, the user located, without remembering the name or any content of the nature tour,

the characteristic vertical blue bar of the web page in the overview images and then used the file list to quickly open the tour's website. A similar event occurred when users were looking for the previously found airfares. Users scrolled through the overview images until seeing the characteristic blue background of Orbitz, and then simply opened the screenshot to view the airfares found in the previous session.

The subjects used the WorkspaceNavigator in the following unanticipated manner. The basic iROS software does not provide a single-step way to send local files from one computer to another. However, the WorkspaceNavigator tools, as a byproduct of capturing a log of the work done, automatically copies all files used to a common directory accessible by all computers in the room. Thus, we observed users opening a file on their machine for the sole purpose of having it go into the WorkspaceNavigator's log so that another user could open the same file on their computer through the information interface. This suggests that an explicit, lightweight mechanism for consolidating information from a collection of machines in an iRoom might be a useful component.

In the second session, we observed an unexpected use of the screenshots when the users realized that they had not saved the Excel file with their budget during the previous session. They simply flipped through the overview images until they saw their spreadsheet on a screen, opened the screenshot to a full size view, and copied the material into a new Excel file. While it may seem redundant to save both visual and active versions of material, this example illustrates that there are cases when the apparent redundancy is useful. Such disjointedness between visual and active information also occurs when the content of a web page changes, and thus the URL displays different information than it did when previously browsed. In the Orbitz case, airfare information found in a previous session was gathered from the screenshot, since it could have been different in the current URL.

In the questionnaires, users rated the usefulness of a variety of features from both the iRoom and the WorkspaceNavigator. As expected, the basic iROS software such as PointRight and MultiBrowse were rated most highly, with WorkspaceNavigator tools, which are used only when needed rather than all the time, rated as somewhat less useful. There was substantial disagreement between users over which features of the WorkspaceNavigator were useful and which were not. The ratings for navigation by keyword and navigation by annotations were approximately the same on average, but 2 users thought the latter was better, 2 thought the former was better, and 2 gave them the same rating. Although the results are not statistically significant, they suggest that individual users work with different models when navigating through captured material, which needs to be considered when improving the interface.

DISCUSSION

These initial studies suggest that people can and will use the material captured by tools such as the WorkspaceNavigator. Users were enthusiastic about the approach; the participants in the first study requested continued use of the system in future meetings even after we had finished our formal study.

As has been demonstrated by others in previous work, annotation is clearly important. Users will annotate extensively if given the chance, primarily to mark critical moments for future recall. Using the iButtons was felt by many users to be as effective, and is certainly less obtrusive, than typing in full text annotations. We discussed, though did not implement, whispered voice annotations, which some users felt would be their ideal annotation model. How best to annotate, and how personal and situational annotation methods are used is an area for further study.

Screenshot capture appears to be a valuable tool. Browsing through the screenshots was a favorite way to review the meeting, both with and without the information interface we provided. It is clear that our interface for accessing screenshots needs to be improved, as it currently takes several steps to see the screenshots in sufficient detail to be fully legible. One interesting problem with displaying screenshots full-sized is that they are often confused with live windows, causing users to try to interact with them.

Another observation from the studies is the value of lightweight mechanisms for creating common organized directories of files. Users generally bring a disparate collection of materials into interactive workspaces for use during unstructured meetings. We believe there is value in this mechanism as our users went to the directory structure our tools automatically created to access files associated with their meetings rather than to the files they stored on their personal machines.

Our experience with the privacy buttons suggests that users have difficulties using real-time privacy controls, even if they have real concerns about privacy. Instead, they prefer to remove sensitive information later. While this offers a less stringent definition of privacy from a technical sense (once captured, there is no guarantee that information will be purged), from a human perspective, it may give more confidence.

While the idea of sessions is a key part of our theoretical model for organizing group work, we have discovered that users do not easily remember to structure sessions as they go. We have proposed techniques for automating session management. However, we suspect that a model using unstructured capture with good tools for structuring and editing afterwards is the best approach.

Our preliminary studies show the power of a visual/spatial organization for augmenting the browsing of the captured material. We saw cases of subjects browsing the timeslices for a particular visual cue (such as the color of a web page) on a particular board (where they remembered it having been displayed). When the WorkspaceNavigator is used over a longer period in a real project, we expect to see this effect amplified.

Finally, there is clear value in having multiple tools record information of similar content in different manners. For instance, we have tools capturing both visual images (screenshots) and active (files and URLs) forms of information. While this information is linked, it is not necessarily the same, as we observed in the cases regarding the unsaved Excel file and the website with changing information. We found it fascinating that the first group used the screenshots, not the active files that were produced, to summarize activities from the meetings,

perhaps because they seemed a simpler and more reliable representation for their purpose.

FUTURE WORK

Our initial studies were conducted to test basic functionality and refine our concepts. The next planned step is to extend the suite of tools provided by the WorkspaceNavigator as part of a systematic study of mechanical engineering design students and their design process over an entire course in spaces other than the iRoom. We are providing users with digital cameras and scanners to explicitly capture non-electronic information, and tools to automatically capture strokes on regular whiteboards. We are easily able to add additional capture tools, such as the whiteboard capture, to the system given our architecture of synchronized independent tools. This study will test a number of hypotheses on how WorkspaceNavigator can be used, and how its use affects measurable project outcomes. We are particularly interested in looking at the use of visual and spatial cues for recall.

Besides the capture model, our studies suggested several improvements for the interface, including improved affordances for making annotations, some automated assistance for defining sessions, and better ways to present the captured information, especially the captured screenshots. We have also upgraded the cameras used for the captured overview images to improve image quality.

Using the wide-angle overview image to organize information in the interface has proven useful, so we want to expand this capability. Currently, it is necessary to run an initial configuration program to outline the extent of each display on the overview image. However, in some settings, the shared displays are moveable, and in all settings we would like to expand the model to automatically include visible laptop screens. We will also generalize the wide-angle view to accommodate setups with multiple cameras.

We have found that the WorkspaceNavigator tools are useful for both meeting participants and for researchers studying meeting participants. These two groups have somewhat different needs, so we are looking into how we can adapt the system specifically for each group. For example, researchers studying meeting participants often desire another layer of information, which includes links to a video or audio stream so they can better understand why as well as what the users are doing. The current interface could be adapted to identify a point of interest, and then a simple action could play the audio or video for that timeslice. We have already done this informally with our own research video recordings, by finding a moment of interest using the WorkspaceNavigator interface, and then using the timestamps to cue our videotapes. We are in the process of augmenting the design study with digital audio capture. Providing such links would also be of value to groups who need videotapes to capture the human interactions as part of their research.

Finally, our studies demonstrate a common pattern where users create unstructured content, then go back and distill it into a more concise form. This

suggests that simplifying reuse is as valuable as simplifying recall, something we would like to explore further.

CONCLUSION

In summary, we have found some key aspects of value in capturing a variety of representations of information generated during organic, unstructured group sessions in interactive workspaces and organizing it in a visual/spatial manner. Our preliminary studies suggest that in spite of some rough edges on the current tools, our mnemonic of using snapshot capture rather than continuous streams of human actions facilitates reuse of information.

ACKNOWLEDGMENTS

We thank the people who helped in the implementation of the WorkspaceNavigator, including Brian Lee, Meredith Ringel, Tico Ballagas, Ya'ir Aizenman, Shankar Ponnekanti, and Pedram Keyani. We also thank everyone who participated in the studies, especially Satinder Gill and Lara Ionescu for organizing the subjects in the first and second studies respectively. This work was supported by a grant from the Wallenberg Global Learning Network and by a grant to Stanford MediaX by the government of Iizuka City, Japan.

REFERENCES

Abowd, G., Atkinson, C., Brotherton, J., Enqvist, T., Gulley, P., LeMon, J. Investigating the Capture, Integration and Access Problem of Ubiquitous Computing in an Educational Setting." In Proceedings of CHI 1998, pp. 440-447.

Ballagas, R., Ringel, M., Stone, M., Borchers, J. iStuff: A Physical User Interface Toolkit for Ubiquitous Computing Environments. To appear in Proceeding of CHI, 2003.

Chiu, P., Kapuskar, A., Reitmeier, S., Wilcox, L. NoteLook: Taking Notes in Meetings with Digital Video and Ink. In Proceedings of the seventh ACM international conference on Multimedia (Part 1), 1999, pp. 149-158.

Chiu, P., Boreczky, J., Grgensohn, A., Kimber, D. LiteMinutes: An Internet-Based System for Multimedia Meeting Minutes. In Proceedings of the Tenth International Conference on World Wide Web, 2001, pp. 140-149.

Czerwinski, M., Horvitz, E. An Investigation of Memory for Daily Computing Events. People and Computers XVI, Proceedings of HCI 2002, pp. 230-245.

Davis, R., Landay, J. et all. NotePals: Lightweight Note Sharing by the Group, for the Group. In Proceedings of CHI 1999, pp. 338-345

DuVaul, R.W., Clarkson, B., Rentalnd, A. The Memory Glasses: Towards a Wearable Context Aware, Situation-Appropriate Reminder System. People and Computer XIV: Proceedings of HCI 2000, Vol 2, British Computer Society, pp. 71-76.

Fass, A. M., Forlizzi J., Pausch, R. MessyDesk and MessyBoard: Two Designs Inspired By the Goal of Improving Human Memory. In Proceedings of DIS 2002, pp. 303-311.

Johanson, B., Ponnekanti, S. R., Sengupta, C., Fox, A. Multibrowsing: Moving Web Content across Multiple Displays. In Proceedings of UBICOMP 2001, pp. 256-272.

Stanford Technical Report TR2002-04

Johanson, B., Fox, A.. The Event Heap: A Coordination Infrastructure for Interactive Workspaces. In Fourth IEEE Workshop on Mobile Computing Systems and Applications, 2002a, pp. 83-93.

Johanson, B., Fox, A., Winograd, T. The Interactive Workspaces Project: Experiences with Ubiquitous Computing Rooms. *IEEE Pervasive Computing*, 2002b, 1(2): pp. 67-74.

Johanson, B., Hutchins, G., Winograd, T., Stone, M. PointRight: Experience with Flexible Input Redirection in Interactive Workspaces. In Proceedings of UIST 2002c, pp. 227-234.

Klemmer, S. R., Thomsen, M., Phelps-Goodman, E. P., Lee, R., Landay, J. A. Where Do Web Sites Come From? Capturing and Interacting with Design History. In Proceedings of CHI 2002, pp. 1-8.

Lamming, M., Flynn, M. Forget-me-not: Intimate Computing in Support of Human Memory. In Proceedings of FRIEND21, International Symposium on Next Generation Human Interface, 1994.

Mark, G., Haake, G., Streitz N. Hypermedia Use in Group Work: Changing the Product, Process and Strategy. In Proceedings of CSCW 1997, pp. 327-368.

Minneman, S., Harrison, S., Janssen, B., Kurtenbach, G., Moran, T., Smith, I., van Melle, B. A confederation of tools for capturing and accessing collaborative activity. In Proceedings of the third ACM international conference on Multimedia 1995, pp. 523-534.

Moran, T., Palen, L., Harrison, S., Chiu, P., Kimber, D., Minneman, S., van Melle, W., Zellweger, P. I'll Get That Off the Audio: A Case Study of Salvaging Multimedia Meeting Records. In Proceedings of CHI 1997, pp. 202-209.

Pedersen, E. R., McCall K., Moran, T. P., Halasz, F. G. Tivoli: An Electronic Whiteboard for Informal Workgroup Meetings. In Proceedings of CHI 1993, pp. 391-398.

Rekimoto, J. Time-Machine Computing: A Time-centric Approach for the Information Environment. In Proceedings of ACM UIST 1999, pp. 45-54.

Renaud, K. Expediting Rapid Recovery From Interruptions by Providing a Visualization of Application Activity. In Proceedings of OZCHI 2000, pp. 348-355.

Richter, H., Abowd, G., Geyer, W., Fuchs, L., Daijavad, S., Poltrock, S. Integrating Meeting Capture within a Collaborative Team Environment. In Proceedings of UBICOMP 2001, pp. 123-128.

Rhodes, B., Starner T. The Remembrance Agent: A Continuously Running Automated Information Retrieval System. In Proceedings of the First International Conference on The Practical Application of Intelligent Agents and Multi Agent Technology, 1996, pp. 487-495.

Robertson, G., Czerwinski, M., Larson, K., Robbins, D. C., Thiel, D., van Dantzich, M. Data Mountain: Using Spatial Memory for Document Management. In Proceedings of UIST 1998, pp. 153-162.

Whittaker, S., Hyland, P., Wiley, M. Filochart: Handwritten Notes Provide Access to Recorded Conversations. In Proceedings of CHI 1994, pp. 271-277.