

# Experiences in Indexing Meeting Content

Heather Richter

Department of Software and Information Systems  
University of North Carolina at Charlotte  
9201 University City Blvd. Charlotte, NC 20223  
richter@uncc.edu

## ABSTRACT

In the workplace, recording activities and discussions can preserve rich, often informal, details of projects and practices, thus contributing to organizational memory. We believe that useful review of these recordings is a knowledge-based task; users will generally be searching and browsing for specific information contained within activities or discussions. In our research, we have designed and implemented a number of meeting capture and access prototypes for different kinds of discussions, exploring ways to index the multimedia meeting recordings. In this paper, we present a number of approaches to content-based indexing and discuss our experiences with them. The approaches highlight some of the opportunities, and the challenges, in structuring recorded meeting information.

## INTRODUCTION

People participate in a number of experiences that we can now automatically record and preserve for later use. In the workplace, recorded activities have the potential to add rich information to the memory of the organization, such as detailed process and status information, and decision or design rationale. Unlike some of the more personal uses of memory capture, workers are usually not concerned with remembering the experience itself, but remembering specific pieces of knowledge contained within that activity. For example, consider recording product design sessions in order to maintain the record of the design decisions and rationale. A designer may wish to review information relating to a particular decision, which may span multiple discussions with a variety of people.

Additionally, this type of organizational memory may serve a broad audience, many of whom were not participants in the original activity itself, or may not even know the those participants. And information within those sessions may be relevant for years, for example as the product is designed, built, tested, maintained, and modified. Finding and reviewing relevant information may be challenging and time consuming, reducing the usefulness of the record.

Thus, in order to be effective, a system that records workplace activities and discussions needs to go beyond preserving and providing access to those activities as separate replayable experiences. Instead, discussions need to be structured and indexed based on the content within them in order to help users find the information they desire.

This need implies two things. First, designers of such systems need to understand the information needs of the users, what kinds of information are part of the activity, and how that knowledge might be used and reviewed. Based on that understanding, systems then need to provide as many relevant structures, indices, summaries and other ways to access and review meeting information.

The research community is making great progress in detecting and recording contextual information – who or what is present in an experience, when and where the activity is occurring. While that context is important to preserve for workplace activities, context will likely be insufficient to effectively structure and interpret the content of activities. Thus, we need to more fully explore methods for creating content-related structures and indices into recorded workplace discussions and activities.

Our work thus far has focused on capturing workplace discussions. We have designed, prototyped, and evaluated several prototypes, enabling us to explore a variety of ways to index recorded meetings. In this paper, we present a set of approaches to creating content-related indices based on our meeting capture experiences. We first summarize the various prototypes we have created, emphasizing the methods for indexing and accessing the content. We then present the set of approaches to content-based indexing, briefly discussing our experiences with each approach and the challenges they present in designing and implementing collective memory systems in the workplace.

## MEETING CAPTURE SYSTEMS

To more fully explore the use of meeting capture and access systems, we designed and implemented prototypes for several different types of meetings – two in the software development domain and one more generic meeting system. We focused in particular on how users can review captured discussions and thus explored a variety of indexing schemes and evaluated their use in finding information. We first briefly summarize each set of prototypes, and in particular, their indexing methods.

**SAAMPad and SAAMPlayer:** SAAMPad is a prototype application that captures software architecture meetings following the Software Architectural Analysis Method (SAAM) [8]. It supports the creation and manipulation of architectural diagrams during structured discussions. The time-stamped interactions with the architectural

components form indices into the captured audio. Unrelated to the SAAMPad implementation, we also built a prototype access interface called SAAMPlayer around video recordings of a set of real SAAM analysis sessions [5]. Users navigated the video based on the locations of keywords highlighted on a timeline. Those keywords were the architectural components and several other key issues discussed during the meetings.

**TAGGER and TagViewer:** With TAGGER, we explored recording knowledge acquisition sessions for large software development projects [6]. As knowledge acquisition members or requirements engineers interview subject matter experts, the TAGGER system records and transcribes the audio. Users then annotate, or “tagged”, the transcript with various software related concepts. For example, users could mark words or phrases that were important things, actions, and constraints within the domain. Users also highlighted potential objects, their attributes and operations, as well as the conversational flow of issues and resolutions. Those categorized annotations then serve as indices into the transcript and video, which we implemented in the interface prototype TagViewer [4]. Users navigated a sample knowledge acquisition video using the text transcript highlighted with the identified concepts. Users could also view summaries of just the tagged phrases, and use them to move to specific points within the discussion.

**TeamSpace:** TeamSpace is a collaborative workspace for distributed teams, with integrated capture and access of meetings [7]. TeamSpace supports generic meetings, providing the ability to create, annotate, and review meetings with agendas, action items, and presentations. Users interact with the capture software during a meeting to create and check off agenda items and action items, change slides, and write or type notes and annotations. These annotated artifacts and events can be reviewed and used to navigate the audio and video of the meeting. This prototype is the only fully functional system, with a fully integrated and automated system for scheduling, recording, and reviewing meetings. Thus, we did not employ more advanced technologies such as speech recognition or location detection in order to more easily deploy a robust system in multiple locations.

## CONTENT-RELATED INDEXING

A variety of meeting capture prototypes have been built in the past decade, exploring novel ways to record video [3], presentations [1] or other meeting information. The most used systems record organized presentations, similar to a classroom lecture. Slides are a natural organizational structure for these presentations. However, slides are not used in every workplace discussion. Many systems also focus on the important context, such as who was present, and the location and time of the meeting. However, as many activities occur that relate to the same project, issue, or decision, users may require ways to not just find one

meeting, but to navigate and link content from many different kinds of discussions.

Many of the scenarios we have written to describe the envisioned use of our various meeting capture prototypes involve individuals returning to a series of discussions looking for detailed information about a decision or design component. Our long term agenda has been to investigate the use of captured meetings for requirements and design rationale and knowledge traceability. These scenarios and goals are very information rich, relying on the long term preservation of large amounts of discussions. If such goals are to be met, we need to find information-centered ways to structure, browse, and search content.

We have still not achieved our goals of creating systems that successfully address knowledge traceability, rationale or other long term uses. However, in designing our prototypes with these goals in mind, we have begun to investigate ways to index the information in a discussion. Each prototype used different indexing schemes based on what was available and appropriate in the discussion, and based on our desire to explore different possibilities. We evaluated each prototype on how those indices helped support users finding detailed information. In evaluations of both SAAMPlayer and TeamSpace, we asked subjects questions and observed them finding answers to those questions in meeting recordings [5]. In evaluating TagViewer, we asked users to prepare a detailed document based on a recorded knowledge acquisition session and observed how they utilized the indices and video [4]. Based on these experiences, we now present and discuss the set of approaches we have explored towards content-based indexing of recorded discussions.

## Meeting structures or process

A number of workplace discussions follow a specific structure or process that can be used to segment the recording. For example, the software architectural analysis method (SAAM) had three distinct stages with different kinds of discussion occurring in each, and SAAMPad behaved differently during these stages. Even general meetings often have an outline or agenda. TeamSpace allowed users to select and check off agenda items during the meeting that help to indicate which topic was being discussed.

In TeamSpace, the agenda provided a very high level structure, indicating the general topics of discussion. In our evaluation, most subjects began a search for information using the agenda [5]. The meetings under review contained only 3 to 5 agenda items, and subjects used the agenda to move to a point in the meeting an average 4 times, approximately once for every question they were asked to answer. Users stated how much they relied on the agenda, and many requested more detailed agendas with subtopics as a finer-grained structure to aid their search. In other words, even this coarse structure was very useful to almost everyone.

The more time between the activity and the review, the more likely users will need general, high-level structures to serve as an overview and reminder of what occurred in the meeting. This structure will not only help them narrow their search, but may spark other memories of the content they are looking for.

However, one of the drawbacks of using meeting structures is their high-level nature. They will likely only indicate large segments of recorded content. SAAM stages could last multiple hours. And very few meetings recorded with TeamSpace had more than 5 agenda items, even meetings that were several hours in length. This meant that while users were able to narrow down their search space, they often still were left with significant chunks of content to search through in other ways. An added challenge is that automatically detecting agendas, topic changes, or other process changes is technically challenging. Systems may need to rely on human effort, albeit minimal, to indicate such changes, at least in the near future.

### **Artifacts**

Many meetings and workplace activities center around relevant artifacts, such as slides, diagrams, and documents. Capturing and linking these artifacts to the discussion is a vital way to provide content-related structure. Like TeamSpace, many capture systems have taken this approach, particularly in recording presentations and whiteboard ink. The SAAMPad system helped users create a very specific artifact –the software architecture. Knowing when users created or highlighted different components of that architecture gives a fairly clear indication that the component was the topic of discussion. Besides simply being indices, artifacts themselves often summarize the information being discussed, providing valuable content even without additional streams of information. An architectural diagram from SAAMPad or notes taken in TeamSpace were valuable even without listening to the audio of the meeting.

People are already accustomed to browsing and searching a number of artifacts. Additionally, text-based documents are often easier to scan and search than audio and video. Thus, being able to search a document for desired information may more quickly lead to also finding the relevant video, or other meeting information.

One big advantage to using artifacts to structure meeting content is that they can be fairly easily captured using today's technology, such as electronic whiteboards, video analysis of slide changes, screen capture, or logging existing artifact software. However, systems may still need to understand or interpret a document to provide meaningful meeting review, such as distinguishing between slide changes in a PowerPoint presentation and scrolling of a Word document. Users may also want integrated views of various artifacts alongside meeting video and other captured indices. Additionally, more specialized artifacts

such as the software architecture diagram in SAAMPad may still require specialized software to capture and review.

### **Keyword**

Another common method for creating content-based indices is using keywords, utilizing either keyword spotting technology or speech-to-text transcripts. This generic method can be applied to any conversation, using keywords taken from other sources or letting users search for their own desired keywords. The SAAMPlayer interface provided the locations of a small set of important keywords, namely the architectural components and several key design issues, to examine the use of this method. Users looked for clusters of keyword locations to find general areas of interest that they then skimmed through. Users often needed to rely on using several keywords to find intersections of those clusters to more manageably narrow down locations of interest.

The biggest challenge in using keywords is that any word, but especially important or common words, will be spoken many times in a discussion. In SAAMPlayer, the names of various architectural components were mentioned 20 to 100 times in six hours of video. Thus, while using keywords will narrow down a search space, finding information will likely still involve skimming or searching those areas further. The more content, the greater this problem will be. With just one hour of video in our TagViewer evaluation, users rarely took advantage of a general keyword search capability and instead relied on the conversation and information indices we discuss next. One user attempted a search by typing in several variations of the same word, including misspellings, to overcome the recognition errors of the transcript. He did not find what he was looking for and returned to using other indices. We need to research more sophisticated information retrieval techniques to understand how to identify the most relevant or interesting segments of discussion based on keywords.

### **Conversation structures**

Another way to structure recorded discussions is to use the flow of the argument or conversation itself. By this we do not mean who is speaking, but instead on what type of statement is being made and potentially the relationships between statements. One example of such a structure is IBIS and other similar design rationale representations [2]. Those techniques suggest identifying the issues, alternatives, and resolutions as they are discussed. We used a simplified version of this structure in TAGGER, allowing users to annotate issues and resolutions in the transcript. A better example is the Compendium system which allows users to create complex conversational maps that can index into meeting video, and be further transformed and integrated with other project information [9].

These types of conversation or argument structures can be applied to a wide variety of meetings, and may reasonably represent important aspects of the discussion. For example,

the issues and resolutions identified in a knowledge acquisition session using TAGGER can help form an initial requirements list. However, identifying the types and structure of conversation statements will likely require significant human effort. While simply identifying issues and resolutions in TagViewer was not very difficult, creating more complex structures like Compendium requires an expert and dedicated facilitator or minute taker.

### Information structures

Even with all of the above indices, finding detailed information when needed may still be too time-consuming. Other more specific structures that pull out, summarize, or indicate the locations of important pieces of information may be needed to support different uses of recorded discussions. Structures that take advantage of what is known about the type or use of the information being discussed could be very powerful. For example, TAGGER based its structures on how knowledge acquisition information would be used, namely within domain and design models and requirements specifications. Thus, users could tag words as important domain terminology or constraints, or as potential objects and attributes. In our evaluation of TagViewer, users were asked to create a requirements document so we expected that the conversational tags of issues and resolutions would be heavily used. Instead, most users preferred the information-specific tags as they provided more of structure and the details of the information they were looking for.

Most users of TagViewer started a search by viewing the lists of tagged words and phrases. Sometimes, just the words within the lists were a sufficient memory aid. Often, users then moved to view the portion of the transcript containing a particular tagged phrase to review the discussion and replay the video. Our sample knowledge acquisition video in our evaluation had, on average, one index every 20 seconds. Thus, these discussions had the most fine-grained structure of any of our prototypes. Users seemed to have less difficulty finding information than in the more sparsely-indexed recordings of TeamSpace and SAAMPlayer. TagViewer users also performed more meeting-wide browsing, using the lists of tagged words to remind them of potential information they had not covered yet. While our explorations are still preliminary, we believe that tools such as TAGGER and TagViewer would be much more useful and usable than generic tools such as TeamSpace because of the more specific information structures that could be utilized.

Yet, information-specific structures are also the most difficult and expensive way to index recorded content. The appropriate information structures will heavily depend on the meeting and information domain, and how that information is to be used. Thus systems that preserve such information may be more specialized and challenging to build. And, much like the conversational structures,

methods for identifying or summarizing important information will likely require significant amounts of human effort, potentially changing the activity itself. In the case of TAGGER, we tried to strike a balance between a useful structure, and one that could be created without too much added burden on meeting participants. While we can still only speculate that the benefits of a system like TAGGER are worth the added effort, we wish to see more systems that investigate such specific and knowledge-oriented methods of structuring and accessing recorded discussions.

### CONCLUSIONS

In our work on meeting capture systems, we have investigated various approaches to indexing and accessing recorded meeting content. We would like to see more systems that capture information-rich discussions and activities, with more complex and specific means for accessing and using recorded content. The approaches we have presented here are not necessarily complete or orthogonal. We hope that through our experiences we have highlighted some of the opportunities, and challenges in designing and building useful collective memory systems in the workplace. We are continuing to research the use of recorded workplace activities, and focusing in particular on information-specific ways to structure and link recorded content.

### REFERENCES

1. P. Chiu, A. Kapuskar, S. Reitmeier, and L. Wilcox. NoteLook: Taking Notes in Meetings with Digital Video and Ink. *Proc. of ACM Multimedia '99*, 1999.
2. J.E. Conklin, and M. Begeman. "Gibis: A Hypertext Tool for Exploratory Policy Discussion." *ACM Transactions on Office Information Systems*, 6(4), 1988, pp 303-331.
3. R. Cutler, Y. Rui, A. Gupta, JJ Cadiz, I. Tashev, L. He, A. Colburn, Z. Zhang, Z. Liu and S. Silverberg, "Distributed meetings: a meeting capture and broadcasting system," *Proc. ACM Multimedia*, Dec. 2002.
4. H. Richter, C. Miller, G.D. Abowd, and I. Hsi. "An Empirical Investigation of Capture and Access for Software Requirements Activities," *Proc. Graphics Interface*, May 2005, pp 121-128.
5. H. Richter. "Designing and Evaluating Meeting Capture and Access Services," Ph.D. Thesis, Georgia Institute of Technology, April 2005.
6. H. Richter, C. Miller, G.D. Abowd, and H. Funk. "Tagging Knowledge Acquisition To Facilitate Knowledge Traceability," *International Journal on Software Engineering and Knowledge Engineering*, World Scientific, 14(1), February 2004, pp 3-19.
7. H. Richter, W. Geyer, L. Fuchs, S. Daijavad, and S. Poltrock. "Integrating Meeting Capture Within a Collaborative Team Environment," *Proc. Ubicomp 2001*, Sept. 2001, pp123-138.
8. H. Richter, P. Schuchhard, and G.D. Abowd. "Automated capture and retrieval of architectural rationale," Position paper, First IFIP Working Conference on Software Architecture, San Antonio, TX, February 1999.
9. A. Selvin, S. Buckingham Shum, M. Sierhuis, J.E. Conklin, B. Zimmermann, C. Palus, W. Drath, D. Horth, J. Domingue, E. Motta and G. Li. Compendium: Making Meetings into Knowledge Events. *Proc. Knowledge Technologies*, Austin, TX, March 2001.