Personal Audio Loop: Reminders from a PAL

Khai N. Truong and Gregory D. Abowd

College of Computing & GVU Center
Georgia Institute of Technology
Atlanta, Georgia 30332-0280, USA
{khai, abowd}@cc.gatech.edu

Abstract. With frequent changes in the topics of discussion or other interruptions, sometimes it is difficult to resume a previous thread in a conversation, even if it occurred only minutes ago. The challenge in recalling salient points to trigger memory recall can be avoided if we can just simply relisten to parts of the conversation. In this paper, we present the Personal Audio Loop (PAL) application, a near-term audio recording system to support the recovery of interrupted conversations. We present the issues in this application space and share how different designs for this system can be implemented to address these issues.

1 Introduction

Sound is an important medium in our daily lives through which we communicate and share ideas. Unfortunately, sound does not persist naturally and fades away with time. As a result, useful content must be remembered if it is to be applied again at a later time. With a large bandwidth of information that can be shared over this medium, it is difficult for us to devote our full attention to remembering everything. Our inability to perceive and identify all potentialy valuable information means no process for selectively remembering information is foolproof. Hence, too often we will struggle to recall information spoken or heard only minutes ago.

During the course of a conversation, the central topic can drift off in a different direction. Interruptions may also occur in the middle of a conversation causing the main point of the discussion to be forgotten. When we simply want to return to a previous topic, we often find ourselves struggling to remember the gist of that conversation, even though it occurred only minutes ago.

Sometimes, we only successfully recover interrupted discussion points after using a few moments to recollect and reshape the conversation. However, too often, we are unable to remember the salient points needed to trigger memory recall and simply give up. None of these scenarios are ideal current practices. Pauses to think over the past few discussion points can be disruptive to the flow of the conversation. Meanwhile, just moving on without being able to continue an important point means we lose a potentially valuable discussion. During formal discussions (e.g., meetings or judicial proceedings), one option for resolving this problem is to employ a stenographer to keep a record of the proceedings that can be consulted to trace shifts between discussion points. However, in informal settings, resorting to manually taking notes requires more effort than is often desired and can potentially prevent people from fully engaging in the conversation.

One of the services of ubiquitous computing is the support for automated capture of streams of information from our daily experiences and future access of these records [3]. Capture and access applications leverage computers to record information, freeing humans to fully engage in the activity. In this paper, we present an application, known as the Personal Audio Loop that explores the near-term capture and access of audio streams to provide a user with a quick retrospective memory of recent events. This system can assist with interruption recovery and activity resumption by allowing audio content to persist for an additional period of time after it has occurred. Unlike a tape recorder, this service continues to capture audio even when playback of previously recorded information is accessed.

1.1 Overview of the Paper

In this paper, we discuss a suite of near-term audio recording systems to support the recovery of interrupted conversations. We begin with a brief review of previous capture and access work that have explored providing the same near-term reminder service offered by our Personal Audio Loop application. There are many ways to provide the personal audio loop service such that it follows the user around wherever she goes. We present the issues in this application space and share how different designs for this system we implemented can address these issues. We conclude with some discussion of future directions in this area.

2 Related Work

Many systems have been built to capture and access experiences in classrooms [1,4,15], meetings [6,16,18], and other generalized experiences [7,10,19]. In previous work, we reviewed many of these systems and identified five domain-specific aspects of the capture and access application space that help designers tailor specific solutions [22]:

Who are the users during capture and access? It is important to identify how many users are supported and whether captured records are public, private or a mixture.

What is captured and accessed? Which artifacts and streams of information are important to review later on, and what level of fidelity is required between the live experience and the playback of the captured experience?

When does capture and access occur? How much time lag is there between the live experience and the expected time of access? How long do captured artifacts need to persist?

Where does capture and access occur? Does capture or access occur in a well-defined location or set of locations? Is mobility during capture or access important?

How is capture and access performed? Are capture and access services provided by the environment or on devices the users carry? What devices and tools must be instrumented in order to facilitate capture of activity and to provide access to a past experience?

These dimensions help define a general design space for capture and access applications that shows how previous research relates and highlights interesting gaps in the design space that remain to be explored. Some examples of these gaps include capture with user mobility and near-term access of captured data.

Existing reminder systems, such as Forget-Me-Not [11], commotion [13], and CybreMinder [8] maintain a memory of important information until a situation arises when the user would need the information again. These context-aware applications investigate how sensed information (such as time, location, and the people around) can be leveraged to provide automatic notification of relevant messages stored and tagged with matching contextual properties —such as a note telling the user to return a book to a friend the next time she and her friend are colocated.

Some long-term capture and access applications (e.g., the Remembrance Agent [17] and WebMemex [23]) have been designed to examine all of a person's captured information to automatically remind users of related information she has previously seen. As the user looks at a document, through recognizing the local context or keywords, these systems remind the user of other documents viewed in the past that are similar.

Near-term reminder systems, instead, provide a user with a quick review of recent events. Near-term capture and access applications that provide audio reminder services have been previously explored in the context of telephone conversations. Xcapture, originally built to provide a "digital tape loop" of a single office, could also provide short-term auditory memory of telephone conversations (5 to 15 minutes long) [10]. The application itself runs on a workstation which captures the phone line and allows the user to quickly review audio content as well as to mark important snippets in the audio loop to save permanently. In MERL's real-time audio buffering technique, the captured audio can persist for the duration of that phone conversation [7]. During the course of a conversation, a user may tap the phone against the ear to move backwards in the audio and to relisten to any portion of the discussion thus far. Audio processing techniques are applied to speed up the playback of the audio to allow the user to continue to hear the recorded audio and eventually catch up to the live content.

The What-Was-I-Cooking application explores the use of collage displays to show recent activities in the kitchen [20]. Through wizard-of-oz techniques, snapshots of key steps performed by a user in the kitchen are selected and annotated with useful information (such as what ingredients or utensils the user is manipulating in that shot) and then added to the collage (which only displays a few recent snapshots at a time). In the case of a memory lapse or interruption, the user can rely on the collage to remind her of the last few things she did. The Where-Were-We application explored providing near synchronous video stream access in a conference room environment [14]. At any time during the meeting, participants can review scenes from minutes ago; however, this service was limited to a single environment.

Because we never know when or where contents from our recent past are needed, a near-term reminder service that follows us is most ideal. Automated capture and access applications can exist on a single device; however, context-aware sensing mechanisms can also be added to multiple environments to allow for captured content to be integrated in a seamless manner. The Conference Assistant application previously demonstrated support for continuous capture and integration in a setting where the user roams between events at a conference setting [9].

3 The Personal Audio Loop Design Space

We now illustrate how to apply the five design dimensions described in the previous section to better understand the specific design space for this near-term capture and access reminder application. In some cases, issues relevant to some of these design dimensions are defined by the problem domain investigated by the application. In other instances, these factors are variables that can change from one design to the next. As a result, we will step through the design dimensions in an order different from previously described. We will then generalize the application design space to identify issues where there exists flexibility (options) to be explored in different designs.

3.1 The What Dimension

The Personal Audio Loop application continuously records the audio stream surrounding a user, where after some time passes the audio is automatically forgotten. During the course of a conversation, when the user needs to be reminded of previous points in the discussion, she may replay any portion of the last 15 minutes of saved audio. No changes are applied to the actual audio content —allowing the user to listen to things again as it was heard the first time. Techniques for removing pauses in the conversation and other ways of speeding up the audio playback [7] can be added to the system later; however, there is no intent to actually modify the verbal content.

3.2 The When Dimension

While we want (ideally) for the capture of the audio content to occur continuously, the intent of the application is for the audio portion of a conversation to be recorded for a user to quickly access any part of the *recent* past, meaning this system only supports the access to near-term recorded conversations. As a result, after some time, the captured audio should be discarded. At this time, we have chosen to allow audio contents to persist for an additional 15 minutes. As we deploy this system for authentic use, this duration may need to change according to individual user preferences.

3.3 The Where Dimension

Since we want to support the continuous capture of a person's audio stream, we must provide this reminder access service where the user is as well. To better ground this work, we begin our exploration in our current work environment (see Figure 1). Our goal is to investigate how to best support the near-term capture and access of our recent audio content in this space, but not limit this service to just one room. The Personal Audio Loop service will *follow* a user around, allowing her to review her recent audio wherever she is located.

People typically enter and exit this environment through the hallway before the "front desk." Informal conversations are often held in the "social area" while the "conference room" is used for different kinds of meetings. The "FCL Lab" area is where most work occurs. Conversations can start anywhere, but often transition into the conference room when conversations turn formal or the social area for friendly chitchat.

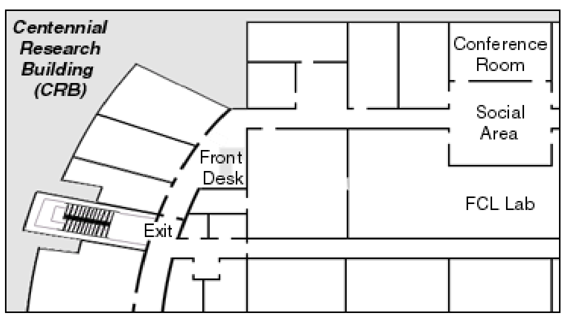

Figure 1. Large office environment with multiple rooms where people converse and the audio reminder service is desired

3.4 The Who Dimension

Each person's audio loop will surely contain different content. The design issue that must be addressed here pertains to whether the application maintains a single audio loop at a time or will it support the capture and access of multiple personal audio loops at once. This question relates directly to *how* the capture and access of the audio loop is supported? On a stand-alone device, the application can be built to simply maintain one person's audio loop. If the environment is instrumented to provide this service, the application must be built to maintain the audio loop for multiple people at once.

3.5 The How Dimension

An automated capture and access application can exist in the simplest form as a single device. As a result, this service can be provided on a device a user carries with her or wears on the body. Alternatively, this service could also be supported through the instrumented physical spaces involved. Figure 2 identifies different possible design options for the Personal Audio Loop application. How to provide this near-term capture and access service can be done in four different ways:

- A person's audio stream is captured and played back on a device she carries with her;
- A person's audio stream is captured and played back by devices instrumented in the environment;
- A person's audio stream is captured by the environment and played back on a device the user carries with her; and
- A person's audio stream is captured on a device the user carries with her and played back by the environment.

In the next section we will present design implementations for each of the four possible configurations. The advantages and disadvantages of each design will be discussed in further detail as well.

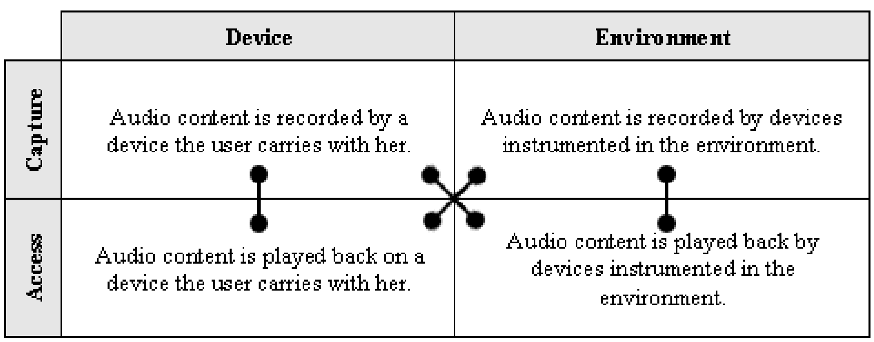

Figure 2. Personal Audio Loop application design space.

4 The Personal Audio Loop Prototypes

We now present a simple instance of the Personal Audio Loop application to better explain how it works. This application was built using the INCA (Infrastructure for Capture and Access) toolkit [21]. INCA provides architectural support for building automated capture and access applications. In order to allow designers to focus on the essential features of a capture and access application, INCA provides a small set of key architectural abstractions and additional features that simplify various aspects of application development of ubiquitous computing systems. We will not discuss INCA itself in this paper, but will explain how the Personal Audio Loop is constructed through a high-level architectural description of how information is captured, stored, and accessed by a user. We will the present how this base application is modified and extended to create prototypes to explore the design space presented in Figure 2.

4.1 The Base Application

In the simplest form, we instrumented a single computer to continuously record the audio portion of a conversation and allow the user to quickly access any part of the recent past. As a Java application, the audio capturer uses the Java Media Framework (JMF) to continuously capture audio in short segments that are timestamped and then stored in a semi-persistent storage. Audio segments are written to disk rather than keeping the audio in the memory buffer. This decision gave us more control of memory usage; however, constant writes to disk to store audio content are CPU intensive. As a result, we chose to segment audio into 4 second chunks to support a near synchronous review of captured information —i.e., almost immediately as soon as something is heard, it can be played back, giving the user time to determine the need to review information and then to formulate and execute the actions to invoke the audio review. After 15 minutes, the captured audio is discarded by the semi-persistent storage. These parameters (4 second chunks and 15 minute persistence) are arbitrary in the sense that each can change depending on the systems requirement of a particular design, as demonstrated later in the paper.

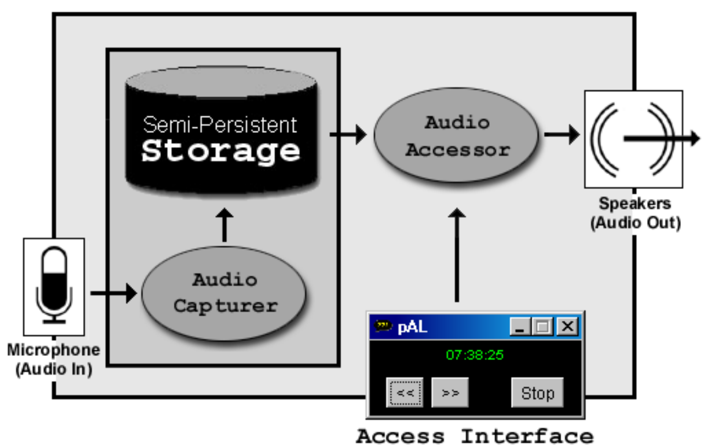

Figure 3. Personal Audio Loop architecture. A near-term audio capture and access reminder system, in which audio from a microphone input is captured and put into a semi-persistent storage (where information lasts for at most 15 minutes). Audio stored in the storage can be accessed by an interface which allows the user to jump back 30 seconds at a time in the audio, 7 seconds forward, or to stop (and return to live).

The access interface (see Figure 3) allows a user to specify a time-point in the last 15 minutes at which to begin review. Audio is then stitched together for playback. The simple user interface was inspired by playback capabilities of a modern digital video recorder, such as TIVO™ and ReplayTV™. When a near-term reminder is desired, the user can jump back 30 seconds at a time into the recently recorded audio stream. She can also nudge forward by 7 seconds at a time in the event of an overshoot. Again, these parameters may change when the system is put into authentic use; however, user testing suggests that these values allow users to quickly move back and forth along the audio loop. More sophisticated skimming techniques are currently under investigation and can be easily introduced into the application.

4.2 A Wearable Interface

Wearable computers have as defining characteristics the features of being portable while operational, proactive, and always on [17]. These properties are ideal for a device that a user can carry with her, always providing her with the last 15 minutes of audio relevant to her.

While output displays are available for wearable computers, they are expensive and one is not necessarily needed for this application. Because the access interface for the base Personal Audio Loop application is fairly simplistic, porting this application to work on a wearable computer required minimal effort. A specialized input device was created to control the playback access of the captured audio. The input device has one button to cause the audio to jump back 30 seconds from its current playback time, and a second button to jump forward by 7 seconds. A third button allows the user to end the playback. Figure 4a shows a wearable computer and the physical interface instrumented to provide the Personal Audio Loop service.

As a wearable unit, a user can carry the unit with her everywhere she goes. As a result, a wearable version of the application can provide the user with the Personal Audio Loop service at all times. However, in actual practice a wearable device has several existing disadvantages as well. While the size of a wearable computer continues to become smaller, users are required to carry along this specialized device. Secondly, the power consumption of a wearable device remains very high. As a result, the device can only continuously provide this service for several hours (3-4 hours) before the batteries need to be recharged. Finally, others with whom a user converses may not know that their audio is being captured. However, because audio only persists for at most 15 minutes, this concern is somewhat alleviated. Despite these disadvantages, a personal device is still ideal in cases where a user wants complete control over when to enable and disable this service. Indeed, we envision a specialized near-term audio appliance, PenPAL, that would fit in a shirt pocket (see Figure 4b).

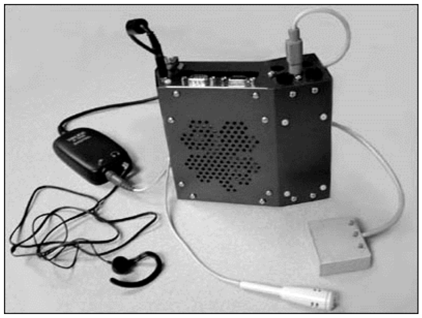

(a)

(b)

Figure 4. Wearable version of the Personal Audio Loop application. 4a) CharmIT™ unit with additional USB audio I/O device (with microphone and earpiece) and specialized input device to control the access captured audio records. 4b) Concept drawing of a specialized device no larger than a pen —PenPAL.

4.3 An Instrumented Environment

Though we found a standalone device providing this service to be useful in our own experience, we wanted to eliminate the need for users to carry any device. The base application was evolved to offload the capture and access functionality to the environment. Recording and playback devices were instrumented in a number of different rooms in our lab environment. To allow users to freely move around in this environment but still have their audio saved, we needed to be able to track each user. Room-level indoor positioning information is used to tag recorded audio describing who was present during each recorded interval. Figure 5 shows the instrumented environment. The end effect was any number of users can move freely in this space, have their audio automatically preserved and be able to access any portion of their recent past anywhere without having to carry around a device.

The positioning sensing system uses a number of floor mats that are antennas for reading RF ID tags. Users who work in this environment each have small RF ID tags which they wear on their shoes or sandals (see Figure 6). Based on which mats a user crosses, the positioning system will identify who all the users are in a room when requested. This system is a simplified version of the work discussed in [2].

Overall, only a few other small changes were made to the original application. Because there existed initial separation of concerns in the base application for capturing and accessing of the audio, we can easily extend the system we have built to handle this distributed situation. The original audio capture component was modified to tag audio segments with user identities and then deployed as a running service in several locations. Each capture service has its own local, semi-persistent storage, which functions the same way as it did before. The access component was modified to allow for a query containing a user identity. All capture and access services are registered and communication is enabled such that audio can be retrieved from different storage in rooms where the user had been previously. Audio segments captured by each room when the user was present will have her name as one of its tagged attributes. These segments are time-ordered, and then audio is stitched together and played back to the user (see Figure 7).

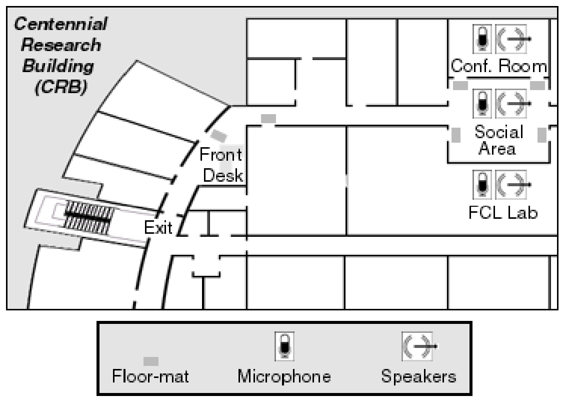

Figure 5. Environment instrumented to capture and tag audio with the name of people present and to support audio playback.

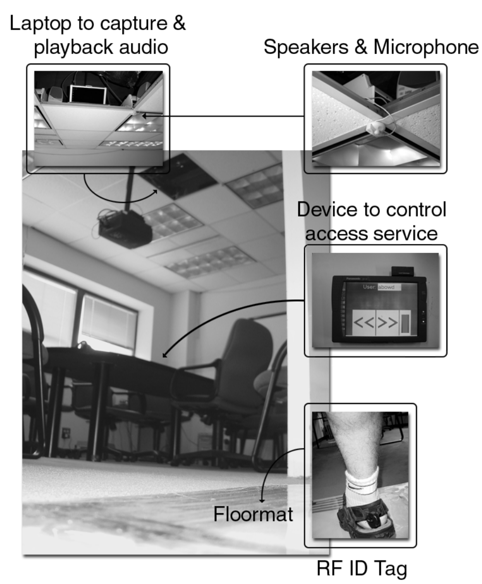

Figure 6. The conference room. Microphones and speakers are placed in the ceiling to record and playback captured audio. Floor-mats are placed at the doors to read and identify users wearing RF ID tags walking over them. Audio content is captured and tagged with the name of people present in that location for the duration of the captured audio segment.

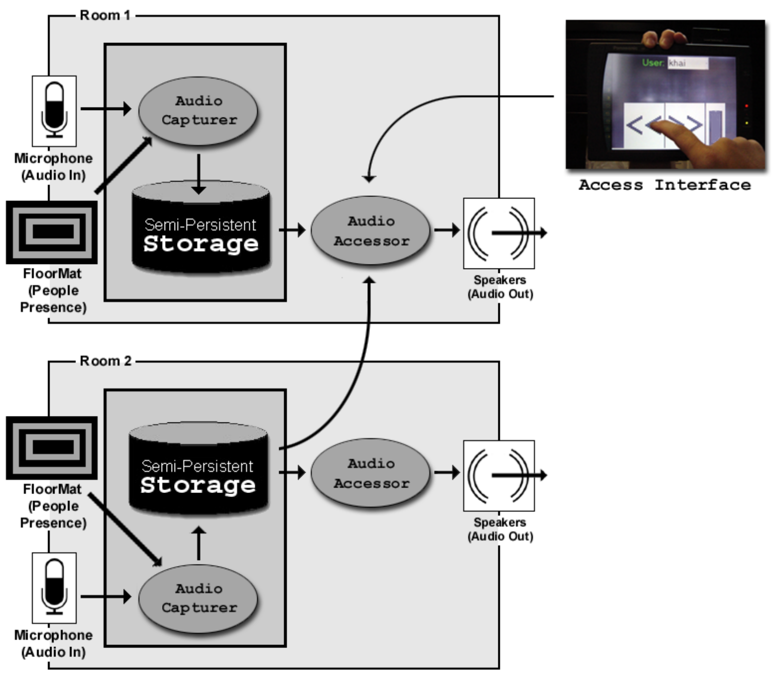

Figure 7. Personal Audio Loop architecture for multiple rooms. Each room is instrumented with microphones, speakers, and floor mats (to identify who are present in the rooms). The Audio Capturer captures audio saved to a semi-persistent storage for that room. Based on where the user is located, the access interface talks to audio accessor for that room, requesting to playback of a portion of that person's audio loop. Audio is then retrieved by the access interface from different storages (depending on where the user has been in the last fifteen minutes), stitched back and played when it is retrieved.

Instead of capturing audio of 4 second segments, this frequency was changed to 1 second instead. When people transition between locations, it is possible for this transition to occur very quickly. When audio is captured in 4 second chunks, the audio stitched back together will not smoothly reflect a user's continuous audio loop. While it is possible to do processing on top of the access and to crop each audio segment to contain only the portions relevant to the user, this means additional processing needed to be performed at play back time. We are interested in making playback as smooth as possible, and wanted to minimize any additional processing that occurs. Meanwhile, changing the capture (segmentation) frequency from 4 seconds to 1 second was a trivial change.

With the environment instrumented to do the capture and playback of each user's personal audio loop, users can expect this service without needing to carry any additional devices with them. Although the cost of instrumenting the environment can be expensive, it is incrementally cheaper for new users to gain this service through the environment rather than needing to have a specialized device (for each user). Additionally, this prototype has required a simple positioning system to be added to the system to make it work. The instrumentation every possible physical space in which a user interacts is clearly not possible. As a result, when users leave the instrumented environment, she will be outside of the coverage of the service. Finally, because public spaces are instrumented to capture audio information, individual privacy is a concern. The application is built on top of INCA, which includes features that allow users to determine the capture and access services available in a physical environment and to control those services, two important features for protecting privacy [12]. However, social protocols must first be studied to identify how to best support the pausing of this reminder service when multiple people are involved.

4.4 Other Permutations

In the ubiquitous computing paradigm, devices are proliferated throughout the environment. A physical space may be instrumented with some useful services but people may also have devices with them. As a result, it is possible to have two other design possibilities:

- A person's audio stream is captured by the environment and played back on a device the user carries with her; and
- A person's audio stream is captured on a device the user carries with her and played back by the environment.

The architecture for an instrumented environment that provides both capture and access services demonstrates how these two other possible designs can be easily implemented as well. Instead of functioning as a stand alone device, the device a user has with her must be networked to talk with the services provided through the environment as well. Our work environment supports a standard 802.11 wireless network, making it easy to network mobile and wearable devices. Currently all devices register with a single *Registry* service running at a well-known location; however, existing research to sense location using the wireless network the devices connect with (such as [5]) could allow us to identify the *Registry* supporting a specific environment when scalability becomes a concern (in cases where users work in multiple instrumented buildings).

4.4.1 Audio is captured by the environment & played back on a personal device

To support access of audio on a personal device when the environment does the capturing, the physical space again must be instrumented to capture audio in the same manner as discussed above. Microphones are placed in the ceilings of each room to record the audio and floormats are used to track users in the environment. Captured audio segments for each room are tagged with the list of people present and then saved for the next 15 minutes. Each user will carry a device that allows her to access

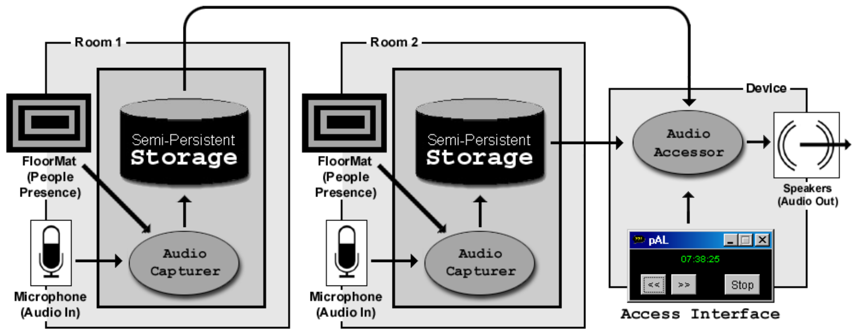

Figure 8. Architecture of a system where capture happens in the environment and access on devices people carry with them. Rooms are instrumented with microphones to record the audio and floormats can provide tags specifying people present in each location. When a user wants to relisten to any portion of her audio loop, audio segments from the semi-persistent storage available in each room are asked for relevant audio segments. Audio is played back as it is retrieved onto the device.

her personal audio loop. Audio segments tagged with her presence are retrieved from the semi-persistent storage of the rooms she was recently in and are stitched together for playback. Figure 8 shows the high-level architecture of this configuration.

Because the device is a personal device, we can assume that it will only want to access the audio of its owner (an attribute easily obtained from the operating system's property settings). Personal access has the advantage of not having others hearing what the user is playing back. Additionally, a device that retrieves audio from the rooms where the user was allows the user to listen to useful captured content even when she steps out of the capture environment.

4.4.2 Audio is captured on a personal device & played back through the environment

It is a natural tendency to share interesting information with people who might not have already heard it. Instead of giving a filtered account of something just heard, a user can carry a personal device that records and saves everything she recently hears. The environment can be instrumented to retrieve the audio from the personal device for playback. As a result, a user can use the playback service to share information with others present in the room. Because captured audio is saved on the personal device, this design gives users control of both when to capture information or not as well as the access of information.

Since audio is captured on a personal device, it can be easily tagged with the user's name. Sound speakers are placed in rooms also containing floormats to playback the audio. When a user specifies that she wants to relisten to a portion of her audio loop, audio accessors instrumented in all the rooms are told of the request. However, only the room where the user is located (as determined by the positioning system using the floormats) will actually retrieve the audio segments from the personal device for playback. Figure 9 shows the high-level architecture of this design configuration.

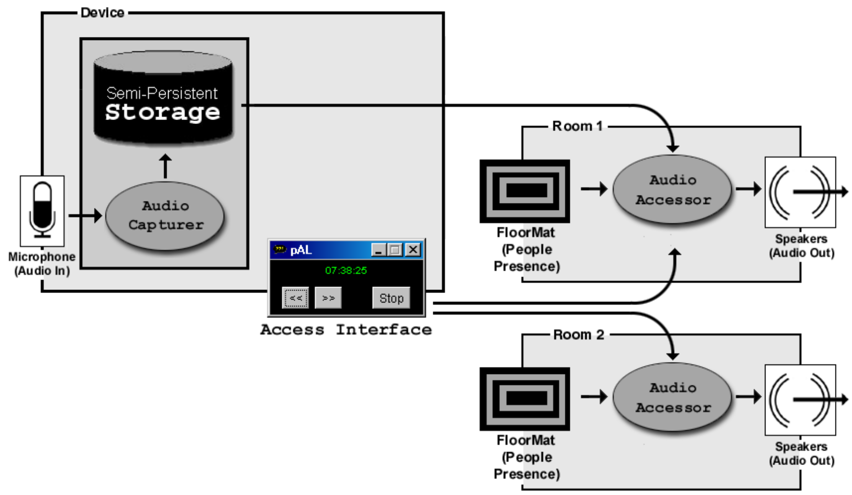

Figure 9. Architecture of a system where personal devices capture a user's audio that is played back in the environment. When playback is desired, audio accessors instrumented in all the rooms are informed of the request; however, only the one where the user is sensed (by the floormat) to be present actually retrieves the audio from the device's semi-persistent storage for playback.

5 Discussion of Research Challenges Explored

As discussed in this paper, the Personal Audio Loop application can be implemented in four different ways. In practice, this service could also exist through both fully instrumented environments and personal devices (i.e., where both support the continuous capture as well as the playback). We did not explicitly discuss this case because it simply a union of the four designs we have presented. As part of this research, we will study which design best serves users in real practice.

Although a stand-alone device can be augmented to effectively capture and playback the audio, if not provided to users in a convenient manner, would a user be willing or remember to carry the device with her? As an option, we have investigated what needs to be added to an environment to achieve the same effect without requiring users to carry anything. To capture personal audio loops for a number of people and still provide a sense of having the service following a user around, rooms are instrumented to tag audio (captured in 1 second chunks) with a list of people present. As a continuous stream of audio can be reconstructed from smaller discrete chunks, relevant audio segments tagged with the user's name are retrieved from various rooms she has been in, and are stitched together in time order for playback upon user request.

When a user wants to relisten to information, it should easy for her to quickly locate the part of the audio she wants to review. The current access interface allows the user to browse the captured audio by moving back 30 seconds at a time and forward by 7 seconds. This uneven browsing method was motivated by existing digital video recorders such as ReplayTV™ and the numbers are adopted after our own initial use of the system. As we deploy the system, we will investigate how to allow the user to efficiently browse the captured audio. We expect these numbers to change according to individual preference. Additionally, users may want different browsing support, such as a visual representation of the audio loop (decorated by silence, location, and speaker identification), or an easy way to jump to points of after pauses or silence in the audio (or complete removal of silences in the audio).

To ease the adoption process, this application is designed to heighten user comfort with the system. While we can see the potential usefulness of saving snippets of audio (such as the feature offered in Xcapture [10]), we chose not to implement this feature at this time. We must first investigate how comfortable people feel having this audio persist for only for a short period of time. As a service that follows the user around, informal conversations are likely to be captured. Whereas previous capture and access applications have investigated the capture and access of information in formal settings (such as the classroom or meeting rooms), informal situations involve people who may not want their audio preserved for posterity. As a result, audio is automatically forgotten after a period of time passes (currently 15 minutes). A second way to improve user comfort is by providing better control of the capture and access services. In instrumented environments, we can provide switches to turn off or pause recording of information. However, conflicts may arise when only some of the people present in a location want the capture to occur. Social protocols for resolving this problem will also need to be investigated.

6 Conclusions & Future Work

Too often in the middle of a conversation an interruption occurs (e.g., the phone rings), and when people try to resume the conversation, they ask the question "What were we talking about?" After a great struggle to recall the last few points of discussion, we simply give up and talk about something different.

The work presented in this paper demonstrates the use of a near-term audio recording system that allows users to playback audio from their recent past. A user can relisten to any portion of the last 15 minutes of saved audio. When discussion threads are forgotten or interruptions occur, this application helps the user recover and resume a previous point in the conversation. There are multiple ways to provide this service to users. Using the five dimensions described in [22], we identified this application design space and demonstrated how different designs can be implemented. In our exploration of this design space, we uncovered many of the technical issues involved in the various ways through which this service can be made available. We discussed these issues and described how we addressed these challenges.

These interfaces will be deployed and we will evaluate the impact of having this enhanced short-term memory available. As a ubiquitous computing application, the Personal Audio Loop will be subjected to real and everyday use. The effects of being able to bring back points in a discussion and how it affects people will be an interesting point of study.

While designed for a work environment, we have outlined a method that can be applied to provide this near-term capture and access service other kinds of environment such as the home. The technique we employed in the Personal Audio Loop system can also be used for different streams of information, such as video. In addition to being able to relisten to things we have recently heard, a Personal Video Loop could allow users to quickly watch again something they just viewed.

Acknowledgements

This work is partially supported by the NSF (Grant 0070345). The authors would like to thank members of the Future Computing Environments Group at Georgia Tech for comments on this work.

References

1. Abowd, G. D. "Classroom 2000: An experiment with the instrumentation of a living educational environment." IBM Systems Journal **38**(4) (1999) 508-530.
2. Abowd, G. D., A. Battestini, et al. The Location Service: A framework for handling multiple location sensing technologies. Submitted to *Ubicomp 2002*, Goteberg, Sweden (2002).
3. Abowd, G. D. and E. D. Mynatt. "Charting Past, Present, and Future Research in Ubiquitous Computing." ACM Transactions on Computer-Human Interaction **7**(1) (2000) 29-58.
4. Bacher, C., R. Muller, et al. Authoring on the Fly. A New Way of Integrating Telepresentation and Courseware Production. In the Proceedings of *ICCE 1997*, Kuching, Sarawak, Malaysia (1997).
5. Castro, P., P. Chiu, et al. A probabilistic room location service for wireless networked environments. In the Proceedings of *Ubicomp 2001*, Atlanta, GA (2001) 18-34.
6. Chiu, P., A. Kapuskar, et al. NoteLook: Taking Notes in Meetings with Digital Video and Ink. In the Proceedings of *ACM Multimedia '99*, Orlando, FL (1999) 149-158.
7. Deitz, P. and W. Yerazunis. Real-Time Audio Buffering for Telephone Applications. In the Proceedings of *UIST'01*, Orlando, Florida (2001).
8. Dey, A. K. and G. D. Abowd. CybreMinder: A Context-Aware System for Supporting Reminders. In the Proceedings of *2nd International Symposium on Handheld and Ubiquitous Computing (HUC2K)*, Bristol, UK (2000) 172-186.

9. Dey, A. K., D. Salber, et al. The Conference Assistant: Combining Context-Awareness with Wearable Computing. In the Proceedings of *ISWC 1999*, San Francisco, CA (1999) 21-28.
10. Hindus, D. and C. Schmandt. Ubiquitous Audio: Capturing Spontaneous Collaboration. In the Proceedings of *CSCW '92*, Toronto, Canada (1992) 210-217.
11. Lamming, M. and M. Flynn. Forget-me-not: intimate computing in support of human memory. In the Proceedings of *FRIEND21 Symposium on Next Generation Human Interfaces*, Tokyo, Japan (1994).
12. Langheinrich, M. Privacy by Design - Principles of Privacy-Aware Ubiquitous Systems. In the Proceedings of *Ubicomp 2001*, Atlanta, GA (2001) 273-291.
13. Marmasse, N. Location-aware information delivery with comMotion. In the Proceedings of *2nd International Symposium on Handheld and Ubiquitous Computing (HUC2K)*, Bristol, UK (2000) 157-171.
14. Minneman, S. and S. Harrison. Where were we: Making and using near-synchronous, pre-narrative video. In the Proceedings of *ACM Multimedia '93* (1993) 1-6.
15. Mukhopadhyay, S. and B. Smith. Passive Capture and Structuring of Lectures. In the Proceedings of *ACM Multimedia '99*, Orlando, FL (1999) 477-487.
16. Pedersen, E. R., K. McCall, et al. Tivoli: An Electronic Whiteboard for Informal Workgroup Meetings. In the Proceedings of *ACM INTERCHI '93*, Amsterdam, The Netherlands (1993) 391-398.
17. Rhodes, B. The Remembrance Agent: A system for augmented memory. In the Proceedings of *1st International Symposium on Wearable Computers (ISWC'97)*, Cambridge, MA (1997) 123-128.
18. Richter, H., G. D. Abowd, et al. Integrating Meeting Capture within a Collaborative Team Environment. In the Proceedings of *Ubicomp 2001*, Atlanta, GA (2001).
19. Stifelman, L. J. The Audio Notebook. Media Laboratory, MIT (1997).
20. Tran, Q. T. and E. D. Mynatt. "What was I Cooking?" An Augmented Kitchen. Submitted to *Ubicomp 2002*, Goteberg, Sweden (2002).
21. Truong, K. N. and G. D. Abowd. INCA: Architectural Support for Building Automated Capture & Access Applications. Submitted to *UIST '02*, Paris, France (2002).
22. Truong, K. N., G. D. Abowd, et al. Who, What, When, Where, How: Design Issues of Capture & Access Applications. In the Proceedings of *Ubicomp 2001*, Atlanta, GA (2001).
23. Truong, K. N., M. Pimentel, et al. Implicit and explicit access of captured media trails. Submitted to *ICME 2002*, Lausanne, Switzerland (2002).